
Fundamental limits for rank-one matrix estimation with groupwise heteroskedasticity

Joshua K. Behne
Duke University

Galen Reeves
Duke University

Abstract

Low-rank matrix recovery problems involving high-dimensional and heterogeneous data appear in applications throughout statistics and machine learning. The contribution of this paper is to establish the fundamental limits of recovery for a broad class of these problems. In particular, we study the problem of estimating a rank-one matrix from Gaussian observations where different blocks of the matrix are observed under different noise levels. In the setting where the number of blocks is fixed while the number of variables tends to infinity, we prove asymptotically exact formulas for the minimum mean-squared error in estimating both the matrix and underlying factors. These results are based on a novel reduction from the low-rank matrix tensor product model (with homogeneous noise) to a rank-one model with heteroskedastic noise.

As an application of our main result, we show that recently proposed methods based on applying principal component analysis (PCA) to weighted combinations of the data are optimal in some settings but sub-optimal in others. We also provide numerical results comparing our asymptotic formulas with the performance of methods based on weighted PCA, gradient descent, and approximate message passing.

1 INTRODUCTION

The problem of extracting information from noisy pairwise data has been studied extensively by focusing on spiked matrix models in which a low-rank signal matrix (the spike) is observed in additive Gaussian noise.

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

The theoretical properties of these models are now well understood in the setting where all of the interactions are of the same type and thus the noise has constant variance. However, much less is known about settings in which different types of information are obtained about different types of pairwise interactions.

Groupwise spiked matrix model This paper provides an asymptotically exact information-theoretic analysis for a general setting in which the underlying objects can be partitioned into different groups and the strength of the pairwise interactions depends on these groups. Specifically, we focus on the K -group spiked matrix model given by

$$Y_{k\ell} = \sqrt{\frac{\lambda_{k\ell}}{N}} \mathbf{x}_k \mathbf{x}_\ell^\top + W_{k\ell}, \quad k, \ell = 1, \dots, K \quad (1)$$

where each \mathbf{x}_k is an unknown $n_k \times 1$ vector of latent random variables and each $W_{k\ell}$ is an unknown $n_k \times n_\ell$ noise matrix with independent standard Gaussian entries. The signal-to-noise ratio in each block is parameterized by a non-negative number $\lambda_{k\ell}$ and a global scale parameter N , which are assumed to be known. Given the collection of observations $Y = (Y_{k\ell})$ and the $K \times K$ matrix $\Lambda = (\lambda_{k\ell})$, the goal is to estimate the latent vectors, (\mathbf{x}_k) , as well as their rank-one outer products, $(\mathbf{x}_k \mathbf{x}_\ell^\top)$.

This model can be expressed in matrix form as

$$Y = G \circ \mathbf{x} \mathbf{x}^\top + W, \quad (2)$$

where $G = (G_{k\ell})$ is a block constant matrix whose $k\ell$ -th block has value $\sqrt{\lambda_{k\ell}/N}$, \circ denotes the Hadamard (entry-wise) matrix product, and \mathbf{x} is column vector of dimension $n_1 + \dots + n_K$ obtained by stacking $\mathbf{x}_1, \dots, \mathbf{x}_K$. We emphasize that G is a known parameter of the model. Hence, if the $\lambda_{k\ell}$ are strictly positive, then each block of the observations can be rescaled to obtain an equivalent model in which the rank-one signal matrix $\mathbf{x} \mathbf{x}^\top$ is corrupted by independent additive Gaussian noise whose variance is constant over blocks.

1.1 Applications

Our theoretical results for the groupwise model in (1) have implications for a variety of applications. In the following, we sometimes use the notation \mathbf{u} and \mathbf{v} instead of \mathbf{x}_k to represent the latent vectors.

PCA with heteroskedastic data Generalizations of the spiked covariance matrix model of Johnstone (2001) have been used to study the effects of heteroskedastic noise (Bai and Yao, 2012; Yao et al., 2015; Hong et al., 2018a,b; Zhang et al., 2021). A rank-one version of such a model can be described as follows. Conditional on a $d \times 1$ latent vector \mathbf{u} , observations are generated independently according to

$$\mathbf{y}_i = \sqrt{\frac{1}{N}} \mathbf{u} v_i + \eta_i \boldsymbol{\epsilon}_i, \quad i = 1, \dots, n \quad (3)$$

where each v_i is a standard Gaussian variable, each $\boldsymbol{\epsilon}_i$ is a $d \times 1$ vector with independent standard Gaussian entries, and the terms N and η_1, \dots, η_L are deterministic parameters.

Hong et al. (2018b) provide asymptotically exact formulas for estimating \mathbf{u} using weighted PCA in the high-dimensional setting where where $d, n \rightarrow \infty$, while the noise levels η_i are supported on fixed set of L values $\{\sigma_1, \dots, \sigma_L\}$. Furthermore, in the setting where the noise levels are known, they describe the optimal choice of weights and the corresponding error rates. Related work by Zhang et al. (2021) focuses on a version of (3) in which the role of \mathbf{u} and $\mathbf{v} = (v_1, \dots, v_n)$ are interchanged and the noise levels are unconstrained.

The model (3) can be seen as a special case of the groupwise model in (1) with $K = L + 1$ where the observations are grouped according to common noise levels such that

$$Y_\ell = \sqrt{\frac{1}{N}} \mathbf{u} \mathbf{v}_\ell^\top + \sigma_\ell W_\ell, \quad \ell = 1, \dots, L \quad (4)$$

where \mathbf{v}_ℓ is an $n_\ell \times 1$ vector with independent standard Gaussian entries, W_ℓ is an $d \times n_\ell$ matrix with independent standard Gaussian entries, and $n_\ell := |\{i : \eta_i = \sigma_\ell\}|$. Here, we have used a slightly different convention for indexing the groups because each observation involves the same factor \mathbf{u} (i.e., the interactions between the \mathbf{v}_ℓ terms are not observed).

As a direct consequence of our main result, we obtain asymptotically exact formulas for the fundamental limits of estimating \mathbf{u} in the setting where the noise levels are known. In the case where \mathbf{u} has IID Gaussian entries, we show that the fundamental limits match the achievable results obtained by Hong et al. (2018b), thus demonstrating that weighted PCA is information-theoretically optimal in this setting; see Theorem 4.

Covariate assisted clustering Another application of interest involves clustering problems with high-dimensional covariates (Binkiewicz et al., 2017; Deshpande et al., 2018; Arroyo et al., 2020). As a specific example, consider the Gaussian version of the contextual stochastic block model introduced by Deshpande et al. (2018), where an unknown random vector $\mathbf{u} \in \{\pm 1\}^d$ has independent Radamacher variables (corresponding to the assignment of each entry to one of two equiprobable communities) and the observations consist of a symmetric $d \times d$ matrix of pairwise observations and an asymmetric $d \times n$ matrix of interactions with an unknown $n \times 1$ latent vector \mathbf{v}

$$Y_{uu} = \sqrt{\frac{\lambda_{uu}}{N}} \mathbf{u} \mathbf{u}^\top + W_{uu} \quad (5a)$$

$$Y_{uv} = \sqrt{\frac{\lambda_{uv}}{N}} \mathbf{u} \mathbf{v}^\top + W_{uv}. \quad (5b)$$

Deshpande et al. (2018) characterize the statistical thresholds for detection and weak recovery of \mathbf{u} and show that these thresholds can be achieved using spectral methods. Ensuing work by Ma and Nandy (2021) proves exact formulas for the asymptotic MMSE and also provides a sharp asymptotic analysis of the MSE associated with an AMP algorithm adapted to this setting.

In view of our groupwise model (1), we see that (5) corresponds to the case of $K = 2$ groups where the pairwise interactions between the entries in \mathbf{v} are not observed (i.e., $\lambda_{vv} = 0$). Applying our main result to this setting recovers the fundamental limits obtained previously as a special case. More generally, our results also apply in setting where λ_{vv} is nonzero and thus one can perform clustering jointly on the entries of \mathbf{u} and \mathbf{v} ; see Section 4.

Computational-to-statistical gaps The analysis of fundamental limits also plays an important role in understanding the interplay between estimation and computational complexity. In this paper we identify two settings where efficient methods are optimal. Theorem 4 gives conditions for the optimality of weighted principal component analysis (WPCA) and our numerical results demonstrate a close correspondence between our asymptotic formulas and the empirical performance of various methods.

Beyond these settings, our result can also be used to study settings where there exists a gap between the optimal performance and efficient methods. Indeed, for a large family of related estimation problems, it is well known that the same formulas describing the optimal limits provide information about the behavior of approximate message passing (Zdeborová and Krzakala, 2016). We leave this as a direction for future research.

1.2 Comparison with Prior Work

Low-rank matrix estimation problems have been studied extensively with applications in covariance estimation Johnstone (2001); Johnstone and Lu (2009), sparse principal component analysis (PCA) Zou et al. (2006), clustering Lesieur et al. (2016); Banks et al. (2018), and community detection Decelle et al. (2011); Abbe (2018). The bulk of the theoretical work has focused on one of two basic models: the *spiked Wigner model*, which corresponds to a single diagonal block of the form Y_{kk} in the groupwise model¹, and the *spiked Wishart model*, which corresponds to a single off-diagonal block of the form Y_{kl} in the groupwise model.

The statistical limits for these models were studied by Lesieur et al. (2015, 2017) who used the replica method to derive conjectured formulas for the asymptotic mutual information and MMSE in the setting of independent and identically distributed variables. These formulas were proven rigorously in ensuing work Krzakala et al. (2016); Barbier et al. (2016); Lelarge and Miolane (2019); Miolane (2017) using a variety of interpolation methods from statistical physics. A complementary line of work has focused on the algorithmic limits of recovery associated with spectral methods Baik et al. (2005); Benaych-Georges and Nadakuditi (2012); Deshpande and Montanari (2014b); Perry et al. (2018), semidefinite relaxations Amini and Wainwright (2009), and approximate message passing (AMP) Fletcher and Rangan (2018); Deshpande and Montanari (2014a); Parker et al. (2014); Lesieur et al. (2017); Montanari and Venkataramanan (2021).

More recent work has focused on extending these models in a number of directions, including models with structured priors described by generative models Aubin et al. (2021); Cocola et al. (2020); Pandit et al. (2020), multiview models involving multiple different observations on the same underlying variables Mayya and Reeves (2019); Barbier and Reeves (2020); Reeves (2020), and low-rank tensor observation problems Luneau et al. (2020); Chen et al. (2021). Within this body of work, there has also been significant interest on the existence of computational-to-statistical gaps under various detection and recovery criteria Barbier et al. (2020); Niles-Weed and Zadik (2020).

In a slightly different context, a special case of the groupwise model (involving noise levels with certain symmetries) has appeared previously as a proof technique Barbier et al. (2018) where the goal was to establish an asymptotic equivalence between the standard

¹The spiked Wigner model assumes a symmetric noise matrix with variance 2 on the diagonal and 1 on the off-diagonal. This corresponds to an independent noise matrix with a factor two difference in SNR.

spiked Wigner model and a *spatially coupled* collection of rank-one models, whose fundamental performance could be analyzed rigorously via the analysis of AMP.

Our theoretical results build upon a long line of work originally in the statistical physics literature including the adaptive interpolation method (Barbier and Macris, 2019) and its recent generalizations of the usual low-rank matrix estimation problems (Barbier and Reeves, 2020; Reeves, 2020). A key step in our analysis is a novel reduction between these models that does not seem to have been recognized previously.

2 FUNDAMENTAL LIMITS

Our main results are exact formulas for the asymptotic mutual information and MMSE associated with the groupwise model (1). In addition to the pairwise observations, we also allow for side information of the form

$$Y_k = \sqrt{r_k} \mathbf{x}_k + W_k, \quad k = 1, \dots, K \quad (6)$$

where each W_k is an unknown $n_k \times 1$ noise vector with independent standard Gaussian entries. The entire collection of observations is denoted by $Y = (Y_k, Y_{k\ell})$.

Assumption 1. We consider a sequence of problems, indexed by an integer N , with observations from the groupwise spiked model (1) and the side information model (6) such that:

- A) The number of groups K is fixed while number of variables in each group $n_k = n_k(N)$ scales as $n_k/N \rightarrow \beta_k \in (0, \infty)$.
- B) For each $k = 1, \dots, K$, the entries of $\mathbf{x}_k = (x_{k1}, \dots, x_{kn_k})$ are drawn independently from a distribution P_k on \mathbb{R} with second moment equal to one and finite fourth moment.

Rather than working with mutual information directly, we use an equivalent parameterization in terms of the relative entropy between the distribution of the observations and the standard Gaussian measure of the same dimension. For each group k , we define the single-letter relative entropy function $D_k: [0, \infty) \rightarrow [0, \infty)$ according to

$$D_k(\gamma) := D(\mathbb{P}_{\sqrt{\gamma} x_k + w_k} \parallel \mathbb{P}_{w_k}) \quad (7)$$

where $x_k \sim P_k$ and $w_k \sim \mathcal{N}(0, 1)$ are independent (univariate) random variables and $D(\mathbb{P} \parallel \mathbb{Q}) = \int \log(\frac{d\mathbb{P}}{d\mathbb{Q}}) d\mathbb{P}$ is the relative entropy (or Kullback-Leibler divergence) between probability measures \mathbb{P} and \mathbb{Q} . Associated with the entire collection of observations, we define $\mathcal{D}_N: [0, \infty)^K \times [0, \infty)^{K \times K} \rightarrow [0, \infty)$ to be

$$\mathcal{D}_N(r, \Lambda) := \frac{1}{N} D(\mathbb{P}_Y \parallel \mathbb{P}_W). \quad (8)$$

Our first result shows that the relative entropy \mathcal{D}_N converges to a well-defined limit that depends only on the parameters (β_k) and the single-letter entropy functions associated with each group. For vectors $u, v \in \mathbb{R}^n$, we write $u \leq v$ if $u_i \leq v_i$ for $i = 1, \dots, n$.

Theorem 1. *Under Assumption 1, the relative entropy function \mathcal{D}_N converges pointwise to the limit \mathcal{D} given by*

$$\mathcal{D}(r, \Lambda) := \max_q \inf_{\tilde{r}} \left\{ \sum_{k=1}^K \beta_k D_k(r_k + \tilde{r}_k) + \frac{1}{2} q^\top \Lambda q - \frac{1}{2} \tilde{r}^\top q \right\} \quad (9)$$

where the maximum is over all K -dimensional vectors q such that $0 \leq q \leq \beta$ and the infimum is over all K -dimensional non-negative vectors \tilde{r} .

The variational formula for the limit provides an explicit link between the prior information about the variables, encapsulated by the single-letter relative entropy functions, and the strengths of the different types of interactions, described by the matrix Λ .

From the point of view of estimation, the significance of Theorem 1 is that changes in the relative entropy with respect to (r, Λ) can be related to the MMSE in estimating the underlying variables. Specifically, as a consequence of the I-MMSE relation (Guo et al., 2005), one finds that \mathcal{D}_N is differentiable on the interior of its domain and its partial derivatives satisfy

$$\begin{aligned} \partial_{r_k} \mathcal{D}_N(r, \Lambda) &= \frac{1}{2N} \mathbb{E}[\|\mathbb{E}[\mathbf{x}_k | Y]\|_2^2] \\ \partial_{\lambda_{k\ell}} \mathcal{D}_N(r, \Lambda) &= \frac{1}{2N^2} \mathbb{E}[\|\mathbb{E}[\mathbf{x}_k \mathbf{x}_\ell^\top | Y]\|_F^2]. \end{aligned} \quad (10)$$

In conjunction with Theorem 1, these relations lead to exact formulas for the asymptotic MMSE. The main difference in the asymptotic setting is that there may be a countable number of points at which the limit is non-differentiable. These points correspond to jump-discontinuities in the asymptotic MMSE.

Theorem 2. *Consider Assumption 1 and suppose that for a given pair (r, Λ) the maximum in (9) is attained at a unique point q^* . For each k such that $q_k^* = 0$ or $r_k > 0$ the MMSE satisfies*

$$\lim_{N \rightarrow \infty} \frac{\mathbb{E}\|\mathbf{x}_k - \mathbb{E}[\mathbf{x}_k | Y]\|_2^2}{n_k} = 1 - \frac{q_k^*}{\beta_k},$$

and for each pair (k, ℓ) such that $q_k^* q_\ell^* = 0$ or $\lambda_{k\ell} + \lambda_{\ell k} > 0$, the MMSE of $\mathbf{x}_k \mathbf{x}_\ell^\top$ satisfies

$$\lim_{N \rightarrow \infty} \frac{\mathbb{E}\|\mathbf{x}_k \mathbf{x}_\ell^\top - \mathbb{E}[\mathbf{x}_k \mathbf{x}_\ell^\top | Y]\|_F^2}{n_k n_\ell} = 1 - \left(\frac{q_k^*}{\beta_k}\right) \left(\frac{q_\ell^*}{\beta_\ell}\right).$$

We emphasize that whether the maximum in (9) is attained at unique point is a property that can be verified on a case by case basis. Furthermore, in Supplement A.4 it is shown that for any fixed r (possibly zero) and fixed collection of off-diagonal terms $(\lambda_{k\ell} : k \neq \ell)$ the set of diagonal terms $(\lambda_{kk} : k = 1, \dots, K)$ such that the maximizer is not unique has Lebesgue measure zero. More generally, there may be cases where some of the limits holds even though the maximizer is not unique.

The reason we need further assumptions on (q^*, r, Λ) is related to the fact that a phase transition with respect to arbitrarily small changes in (r, Λ) occurs when the global maximizer jumps from one location to another. For a more thorough discussion of this behavior see (Reeves, 2020, Section II-B). While the conditions we provide are sufficient, they are probably not necessary.

Finally, there may be settings of interest where r is zero and Λ is sparse for which the assumptions of Theorem 2 are not satisfied. The next result shows that every global maximizer provides a lower bound on the MMSE. This result is used in Section 3 to establish the optimality of weighted PCA.

Theorem 3. *Consider Assumption 1 and suppose that for a given pair (r, Λ) the maximum in (9) is attained at q^* . Then, the MMSE satisfies*

$$\begin{aligned} \liminf_{N \rightarrow \infty} \frac{\mathbb{E}\|\mathbf{x}_k - \mathbb{E}[\mathbf{x}_k | Y]\|_2^2}{n_k} &\geq 1 - \frac{q_k^*}{\beta_k}, \\ \liminf_{N \rightarrow \infty} \frac{\mathbb{E}\|\mathbf{x}_k \mathbf{x}_\ell^\top - \mathbb{E}[\mathbf{x}_k \mathbf{x}_\ell^\top | Y]\|_F^2}{n_k n_\ell} &\geq 1 - \left(\frac{q_k^*}{\beta_k}\right) \left(\frac{q_\ell^*}{\beta_\ell}\right). \end{aligned}$$

To help interpret these results it is useful to consider some special cases. Observe that if there is no side information ($r = 0$) and Λ is diagonal, then the groupwise model decouples into K independent spiked Wigner models. In this case, the limit can be expressed as

$$\sum_{k=1}^K \max_{q_k \in [0, \beta_k]} \left\{ \beta_k D_k(2\lambda_{kk} q_k) - \frac{1}{2} \lambda_{kk} q_k^2 \right\}. \quad (11)$$

Here, each summand corresponds to the limits obtained previously for the spiked Wigner model Barbier et al. (2016) with the only difference being a factor of 2 in the signal-to-noise ratio which arises because our model does not assume symmetric noise.

Alternatively, if there is no side information, Λ is anti-diagonal, and K is even, then the groupwise model decouples into $K/2$ independent Wishart models. After some straightforward manipulations, the limit can be expressed as

$$\sum_{k < \ell : k + \ell = K} \max_{q_k \in [0, \beta_k]} \min_{q_\ell \in [0, \beta_\ell]} \left\{ \beta_k D_k(2\bar{\lambda}_{k\ell} q_\ell) + \beta_\ell D_\ell(2\bar{\lambda}_{k\ell} q_k) - \bar{\lambda}_{k\ell} q_k q_\ell \right\} \quad (12)$$

where $\bar{\lambda}_{k\ell} = (\lambda_{k\ell} + \lambda_{\ell k})/2$. Here, each summand corresponds to the limit obtained previously for the spiked Wishart model Miolane (2017).

2.1 Proof Outline

A key step in the proof of Theorem 1 is to recognize that the groupwise model can be expressed as a special case of the matrix tensor product (MTP) model studied by Reeves (2020). The MTP model was conceived as a generalization of multiview low-rank estimation models in which one obtains multiple observations associated with a collection of d -dimensional variables represented by an $n \times d$ matrix X . The observations in this model have the form

$$\tilde{Y} = (X \otimes X)B + \tilde{W} \quad (13)$$

where $X \otimes X$ is the $n^2 \times d^2$ matrix formed by the Kronecker matrix product, B is a known $d^2 \times m$ coupling matrix, and \tilde{W} is an $n^2 \times m$ noise matrix with independent Gaussian entries. The main result in Reeves (2020) is a formula describing the exact limit of the relative entropy in the asymptotic regime where the dimension d is fixed while the number of variables n increases to infinity.

The connection between the groupwise model in (1) and the MTP model can be understood as follows. First, let the latent vectors $\mathbf{x}_1, \dots, \mathbf{x}_K$ be embedded into an $n \times K$ block diagonal matrix of the form $X = \text{diag}(\mathbf{x}_1, \dots, \mathbf{x}_K)$. Then, the matrix version of the groupwise model in (2) can be expressed as

$$Y = \frac{1}{\sqrt{N}} X \Gamma X^\top + W, \quad \Gamma := (\sqrt{\lambda_{k\ell}}). \quad (14)$$

Next we use vectorization to express this matrix equation in terms of the Kronecker product. For an $m \times n$ matrix A , let $\text{vec}(A)$ be the $mn \times 1$ vector obtained by stacking the columns of A . We have

$$\text{vec}(Y) = \frac{1}{\sqrt{N}} (X \otimes X) \text{vec}(\Gamma) + \text{vec}(W). \quad (15)$$

Hence, our model corresponds to the matrix tensor product where the dimension is equal to the number of groups and the coupling ‘‘matrix’’ is described by the $k^2 \times 1$ vector $\text{vec}(\Gamma)$.

From here, the remaining step in the proof is to simplify the expression for the limiting relative entropy. A direct application of the results in Reeves (2020) leads to a formula that requires optimization over the space of $K \times K$ symmetric positive semidefinite matrices (see (Reeves, 2020, Definition 3)). Exploiting the block-diagonal structure of the matrix X arising from the groupwise model, we are able to simplify this expression to obtain the formula given in Theorem 1, in which the

optimization is taken over the space of K -dimensional nonnegative vectors.

The proof of Theorem 2 follows from Theorem 1 using a similar approach as the proof of (Reeves, 2020, Theorem 2). The full proofs of Theorems 1 and 2 are provided in Supplement A.2 and A.3 respectively.

3 OPTIMALITY OF WEIGHTED PCA

This section considers some implications of our main results for the heteroskedastic model given in (4). Recall that in this setting the observation consists of interactions between a latent vector \mathbf{u} and a collection Gaussian vectors (\mathbf{v}_ℓ) , where the noise level in each interaction is parameterized by a known value σ_ℓ . Throughout this section we further assume that d, n_1, \dots, n_L scale to infinity with the global scale parameter N such that $d/N \rightarrow \beta_0 \in (0, \infty)$ and $n_\ell/N \rightarrow \beta_\ell \in (0, \infty)$.

A standard approach for estimating \mathbf{u} from the observations Y_1, \dots, Y_L is to use sample weighted principle component analysis (WPCA), where the estimate $\hat{\mathbf{u}}$ is a scaled version of the leading eigenvector of the $d \times d$ positive semidefinite matrix

$$\sum_{\ell} \omega_{\ell} Y_{\ell} Y_{\ell}^{\top} \quad (16)$$

Hong et al. (2018b) study the asymptotic performance of WPCA for a low-rank version of (4) where \mathbf{u} is replaced by a matrix with orthogonal columns. For the rank-one setting considered in this paper, their results provide formulas for the squared correlation between $\hat{\mathbf{u}}$ and \mathbf{u} as a function of the parameters $(\sigma_{\ell}, \beta_{\ell})$ as well as the weights (ω_{ℓ}) . Furthermore, they derive the optimal choice of weights as a function of the problem parameters and provide a simplified formula for the corresponding correlation. Adapted to the notation of this paper, their results show that

$$\frac{\langle \hat{\mathbf{u}}, \mathbf{u} \rangle^2}{\|\mathbf{u}\|^2 \|\hat{\mathbf{u}}\|^2} \xrightarrow{a.s.} q_0^{\text{WPCA}} / \beta_0 \quad (17)$$

where $\hat{\mathbf{u}}$ is obtained via the optimal choice of weights and q_0^{WPCA} is given by the largest real root of the function

$$R(x) := 1 - \sum_{\ell=1}^L \frac{\beta_{\ell} \beta_0 - x}{\sigma_{\ell}^2 \sigma_{\ell}^2 + x} \quad (18)$$

To compare results for WPCA with the formulas for the MMSE obtained in this paper, observe that the model in (4) can be represented as an instance of the groupwise model (1) with $K = L + 1$ groups indexed

by the set $\{0, 1, \dots, L\}$ via the mapping

$$\mathbf{u} = \mathbf{x}_0, \quad \mathbf{v}_\ell = \mathbf{x}_\ell, \quad \lambda_{k\ell} = \begin{cases} \sigma_\ell^{-2}, & k = 0, \ell \geq 1 \\ 0, & k \geq 1. \end{cases}$$

Notice that in this case, the quadratic term in the variational formula (9) simplifies as

$$q^\top \Lambda q = q_0 \sum_{\ell=1}^L q_\ell / \sigma_\ell^2 \quad (19)$$

Consequently, for $r = 0$ the expression (9) is given by

$$\max_{0 \leq q \leq \beta} \inf_{\tilde{r} \geq 0} \left\{ \frac{1}{2} \sum_{k=0}^L \beta_k D_k(\tilde{r}_k) + \frac{q_0}{2} \sum_{\ell=1}^L \frac{q_\ell}{\sigma_\ell^2} - \frac{1}{2} \tilde{r}^\top q \right\} \quad (20)$$

Theorem 4. *Consider Assumption 1 and suppose that \mathbf{u} and $\mathbf{v}_1, \dots, \mathbf{v}_L$ have independent standard Gaussian entries. Then, (20) has a unique maximizer q^* . Furthermore, if*

$$\sum_{\ell=1}^L \frac{\beta_0 \beta_\ell}{\sigma_\ell^4} > 1 \quad (21)$$

then $q_0^* = q_0^{\text{WPCA}}$. Otherwise $q_0^* = 0$.

Theorem 4 shows that WPCA attains the fundamental limits described by our formulas. Moreover, it gives an explicit condition on the noise levels such that the correlation is strictly positive. The full proof of Theorem 4 is given in Supplement A.5.

4 NUMERICAL RESULTS

To simplify the presentation of the algorithmic results, it is convenient to apply a preprocessing step that symmetrizes the observations. The symmetrized observations and parameters are defined according to

$$Y_{k\ell}^{\text{sym}} = \sqrt{\frac{\lambda_{k\ell}}{\lambda_{k\ell} + \lambda_{\ell k}}} Y_{k\ell} + \sqrt{\frac{\lambda_{\ell k}}{\lambda_{k\ell} + \lambda_{\ell k}}} Y_{\ell k}^\top \quad (22)$$

$$\lambda_{k\ell}^{\text{sym}} = \lambda_{k\ell} + \lambda_{\ell k},$$

where we use the convention $Y_{k\ell}^{\text{sym}} = \sqrt{1/2} Y_{k\ell} + \sqrt{1/2} Y_{\ell k}^\top$ in the case $\lambda_{k\ell}^{\text{sym}} = 0$. By the orthogonal invariance of the standard Gaussian distribution, it can be verified that $(Y_{k\ell}^{\text{sym}})$ are sufficient statistics for estimation of (\mathbf{x}_k) . The symmetric groupwise model can be written in matrix form according to

$$Y^{\text{sym}} = G^{\text{sym}} \circ \mathbf{x}\mathbf{x}^\top + W^{\text{sym}}, \quad (23)$$

where \mathbf{x} is the $n \times 1$ vector obtained from stacking the (\mathbf{x}_k) , W^{sym} is an $n \times n$ symmetric matrix drawn from the Gaussian orthogonal ensemble, and G^{sym} is an $n \times n$ block-constant matrix with entries in the $k\ell$ -th block given by $N^{-1/2} (\lambda_{k\ell}^{\text{sym}})^{1/2} = N^{-1/2} (\lambda_{k\ell} + \lambda_{\ell k})^{1/2}$.

4.1 Spectral Methods

Spectral methods have been studied extensively both in the context of the spiked Wigner and spiked Wishart models, as well as covariate assisted models and multiview models. In the case of single-type interactions, the spectral methods are optimal when the factors have an orthogonally invariant distribution. We consider two baseline methods based on principal component analysis (PCA):

- **Joint PCA:** Let Y^* be the symmetric $n \times n$ data matrix obtained by zeroing out the blocks in Y^{sym} with $\lambda_{k\ell}^{\text{sym}} = 0$. The estimate of each \mathbf{x}_k is chosen to be proportional to the corresponding sub-vector in the leading eigenvector of Y^* .
- **Weighted PCA:** Given nonnegative weights $(w_{k\ell})$ the estimate of each \mathbf{x}_k is chosen to be proportional to the leading eigenvector of the $n_k \times n_k$ matrix

$$w_{kk} Y_{kk}^{\text{sym}} + \sum_{\ell \neq k} w_{k\ell} Y_{k\ell}^{\text{sym}} (Y_{k\ell}^{\text{sym}})^\top. \quad (24)$$

In the case of the two-group model, this method is equivalent to the one proposed by Deshpande et al. (2018, Theorem 6), who provide an explicit specification for the weights w_{12} and w_{21} as a function of the noise level in each block. Alternatively, if the Y_{kk} observation is omitted (i.e., $w_{kk} = 0$) then this reduces to the method studied in Section 3.

4.2 Gradient Descent

Gradient descent algorithms have been studied in many contexts and provide a straightforward optimization procedure. Recent work (e.g., Zhang and Zhang (2020); Bhojanapalli et al. (2016)) has shown that, under reasonable conditions, low-rank matrix recovery has an objective without any spurious local minima, even though it is a non-convex problem. A gradient method was also studied in the spiked matrix model under generative priors Cocola et al. (2020).

We approximate the maximum likelihood estimate of the symmetric model (23) with gradient descent. Let \mathbf{x}^0 be an initial estimate of \mathbf{x} and let $\gamma_0, \gamma_1, \dots$ be a sequence of positive numbers. For $t = 0, 1, 2, \dots$ the gradient descent updates are given by

$$\mathbf{x}^{t+1} = \mathbf{x}^t + \gamma_t \left[(G^{\text{sym}} \circ Y^{\text{sym}}) \mathbf{x}^t - ((G^{\text{sym}})^{\circ 2} (\mathbf{x}^t)^{\circ 2}) \circ \mathbf{x}^t \right] \quad (25)$$

where $(A)^{\circ 2}$ denotes the entry-wise square for an arbitrary matrix or vector A . The updates can be expressed

equivalently at the group level according to

$$\mathbf{x}_k^{t+1} = \mathbf{x}_k^t + \gamma_t \left[\sum_{\ell=1}^K \sqrt{\frac{\lambda_{k\ell}^{\text{sym}}}{N}} Y_{k\ell}^{\text{sym}} \mathbf{x}_\ell^t - \left(\frac{1}{N} \sum_{\ell=1}^K \lambda_{k\ell}^{\text{sym}} \|\mathbf{x}_\ell^t\|^2 \right) \circ \mathbf{x}_k^t \right] \quad (26)$$

where $k = 1, \dots, K$.

4.3 Approximate Message Passing

A variety of AMP algorithms have been proposed in the context of low-rank matrix estimation problems Fletcher and Rangan (2018); Deshpande and Montanari (2014a); Parker et al. (2014); Lesieur et al. (2017); Montanari and Venkataramanan (2021); Ma and Nandy (2021). This section gives a version of AMP adapted to the symmetric form of the groupwise model (23). The derivation follows the general outline given by Lesieur et al. (2017) and is described in full detail in Supplement B.

For $k = 1, \dots, K$, define the function

$$\eta_k(a, b) = \frac{\int x_k \exp\{ax_k - \frac{1}{2}bx_k^2\} dP_k(x_k)}{\int \exp\{ax'_k - \frac{1}{2}b(x'_k)^2\} dP_k(x'_k)}, \quad (27)$$

and let $\eta'(a, b)$ denote the derivative with respect to the first argument. Also, for each $k = 1, \dots, K$, let \mathbf{m}_k^0 be an initial estimate of \mathbf{x}_k . For $t = 0, 1, 2, \dots$ the AMP updates are given by

$$\mathbf{a}_k^t = \sum_{\ell=1}^K \sqrt{\frac{\lambda_{k\ell}^{\text{sym}}}{N}} Y_{k\ell}^{\text{sym}} \mathbf{m}_\ell^t - \left(\frac{1}{N} \sum_{\ell=1}^K \lambda_{k\ell}^{\text{sym}} (Y_{k\ell}^{\text{sym}})^{\circ 2} \mathbf{v}_\ell^t \right) \circ \mathbf{m}_k^{t-1} \quad (28a)$$

$$\mathbf{b}_k^t = \frac{1}{N} \sum_{\ell=1}^K \left(\lambda_{k\ell}^{\text{sym}} (\mathbf{m}_\ell^t)^{\circ 2} + \lambda_{k\ell}^{\text{sym}} (\mathbf{1}_{n_k \times n_\ell} - (Y_{k\ell}^{\text{sym}})^{\circ 2}) \mathbf{v}_\ell^t \right) \quad (28b)$$

$$\mathbf{m}_k^{t+1} = \eta_k(\mathbf{a}_k^t, \mathbf{b}_k^t) \quad (28c)$$

$$\mathbf{v}_k^{t+1} = \eta'_k(\mathbf{a}_k^t, \mathbf{b}_k^t), \quad (28d)$$

where \mathbf{m}_k^{-1} and \mathbf{v}_k^0 are initialized to the all zeros vector, the functions η_k and η'_k are applied entrywise to vector-valued inputs, and all of the terms for $k = 1, \dots, K$ are updated at each iteration. After T iterations, the approximation of the conditional mean \mathbf{x}_k is given by \mathbf{m}_k^T and the approximation to the conditional variance is given by \mathbf{v}_k^T . Furthermore, the vector \mathbf{a}_k^T can be viewed as an unbiased estimate of \mathbf{x}_k . See Supplement B for more details.

4.4 Discussion of Empirical Results

This subsection provides empirical results comparing the formulas for the asymptotic MMSE in Theorem 2 and the empirical performance of the methods described in previous subsections. We will focus our attention on the groupwise model (1) with $K = 2$ groups and equal group sizes. We also parameterize our $\Lambda \in \mathbb{R}^{2 \times 2}$ matrix by $\alpha \in [0, 1]$ and $\lambda \in (0, \infty)$. This yields the following model

$$\Lambda(\alpha, \lambda) = \lambda \begin{pmatrix} 1 - \alpha & \alpha \\ \alpha & 1 - \alpha \end{pmatrix}, \quad \mathbf{x} = \begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix} \quad (29)$$

where $\mathbf{u} \in \mathbb{R}^{n_u}$ and $\mathbf{v} \in \mathbb{R}^{n_v}$ with $n_u = n_v = n$. The MMSE for this model can be obtained from the global optimizer, (q_1^*, q_2^*) , of

$$\max_{q_1, q_2} \inf_{\tilde{r}_1, \tilde{r}_2} \left\{ \frac{1}{2} D_1(\tilde{r}_1) + \frac{1}{2} D_2(\tilde{r}_2) - \frac{1}{2} (\tilde{r}_1 q_1 + \tilde{r}_2 q_2) + \frac{\lambda}{2} [(1 - \alpha)(q_1^2 + q_2^2) + 2\alpha q_1 q_2] \right\}, \quad (30)$$

where D_1 and D_2 are the single-letter relative entropy functions associated with the distributions of \mathbf{u} and \mathbf{v} , respectively, and we have taken the global scale parameter in Theorem 1 to be $N = 2n$.

It can be observed that the fundamental limits of this model are invariant to the parameter α when the priors on \mathbf{u} and \mathbf{v} are the same. When this is the case in our empirical results, we will plot only a single curve for the MMSE (i.e., the curve generated by fixing any value of α and varying λ).

Evaluation of MSE Throughout our empirical results, performance is assessed in terms of the MSEs in estimating the rank-one matrices $\mathbf{u}\mathbf{u}^\top$ and $\mathbf{v}\mathbf{v}^\top$ appearing on the diagonal blocks of $\mathbf{x}\mathbf{x}^\top$. In all cases, the estimates of the matrices are obtained from the estimates $\hat{\mathbf{u}}$ and $\hat{\mathbf{v}}$ of the corresponding vectors and the MSE can be expressed as

$$\begin{aligned} \text{mse}_{uu} &= \frac{1}{n^2} \mathbb{E}[\|\mathbf{u}\mathbf{u}^\top - \hat{\mathbf{u}}\hat{\mathbf{u}}^\top\|_F^2] \\ \text{mse}_{vv} &= \frac{1}{n^2} \mathbb{E}[\|\mathbf{v}\mathbf{v}^\top - \hat{\mathbf{v}}\hat{\mathbf{v}}^\top\|_F^2] \end{aligned} \quad (31)$$

For the gradient descent method and AMP these expectations are approximated directly via Monte Carlo trials. For the spectral methods described in Section 4.1, the estimates of \mathbf{u} and \mathbf{v} are defined up to an unspecified scale factor. For the purposes of this paper, we focus on the best possible MSE that can be achieved under the constant-norm constraint $\|\hat{\mathbf{u}}\| = \rho_u$ and $\|\hat{\mathbf{v}}\| = \rho_v$ where ρ_u and ρ_v are fixed values that are optimized as a function of the distribution of the data.

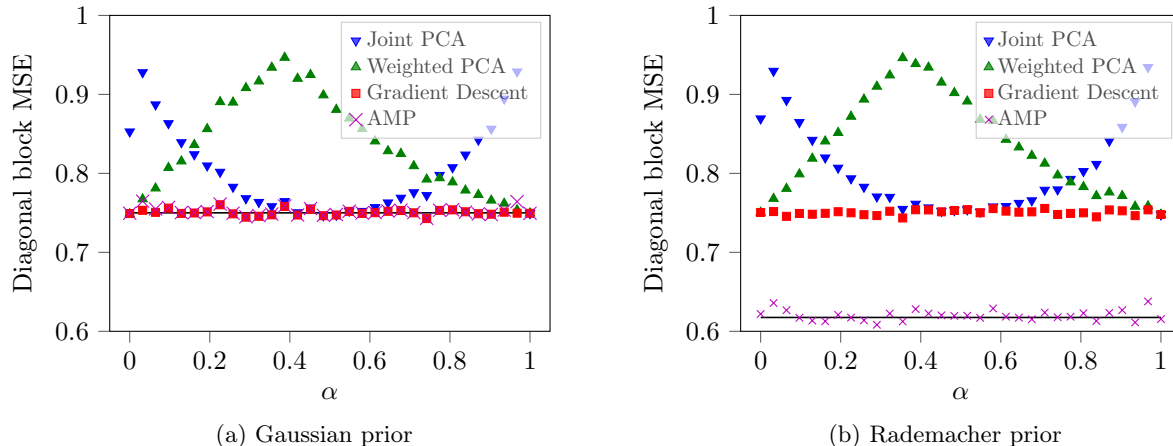


Figure 1: MSE in estimating the rank-one matrices $\mathbf{u}\mathbf{u}^\top$ and $\mathbf{v}\mathbf{v}^\top$ in the two-group model (29) as function of α with $n = 1024$ and $\lambda = 2$. The entries of \mathbf{x} are IID standard Gaussian in (a) and IID Rademacher in (b). The asymptotic MMSE (solid line) is given by Theorem 2. The diagonal block MSE is computed by averaging over 64 Monte Carlo trials for each algorithm.

Starting with an initial estimate $\tilde{\mathbf{u}}$ (assumed to be nonzero with probability one) the minimum MSE that can be obtained by projecting $\tilde{\mathbf{u}}$ onto a sphere of fixed radius can be expressed as

$$\begin{aligned} \inf_{\rho_u \geq 0} \frac{1}{n^2} \mathbb{E} \left[\left\| \mathbf{u}\mathbf{u}^\top - \rho_u^2 \frac{\tilde{\mathbf{u}}\tilde{\mathbf{u}}^\top}{\|\tilde{\mathbf{u}}\|^2} \right\|_F^2 \right] \\ = \frac{1}{n^2} \mathbb{E} [\|\mathbf{u}\|^4] - \frac{1}{n^2} \mathbb{E} \left[\frac{(\mathbf{u}^\top \tilde{\mathbf{u}})^2}{\|\tilde{\mathbf{u}}\|^2} \right]^2. \end{aligned} \quad (32)$$

Using this identity, we obtain an unbiased estimate of the desired MSE by approximating the second expectation on the right-hand side using Monte Carlo trials.

Implementation details The gradient descent method is initialized using estimates obtained from joint PCA, and AMP is initialized with each \mathbf{m}_k^0 initialized with IID $\mathcal{N}(0, 10^{-6})$ entries. For the weighted PCA method, the weights are chosen via a grid search optimization procedure that minimizes the empirical loss. Note that in doing so, the weights are adapted to the ground truth values of \mathbf{u} and \mathbf{v} , and thus the resulting MSE should be viewed as a lower bound on the performance of the weighted PCA method for any data-driven selection of the weights. Additional details regarding the numerical approximation of the MMSE formulas are given in Supplement C. In addition, the code used to produce our numerical results can be found at <https://github.com/joshuabehne/GROME>.

Gaussian priors The case where both \mathbf{u} and \mathbf{v} have independent standard Gaussian entries is shown in Figure 1a. In this setting, the asymptotic MMSE is constant for all $0 \leq \alpha \leq 1$ and the weak detection

threshold (i.e., the smallest value of λ such that the asymptotic MMSE is strictly less than the MMSE obtained without any observations) occurs at the critical threshold $\lambda^* = 1$. The empirical performance of both gradient descent and AMP shows excellent agreement with the formulas for the asymptotic MMSE. By contrast, the spectral methods are sensitive to the choice of the parameter α . Additional plots for scenarios with Gaussian priors are given in Supplement D.

Non-Gaussian priors A comparison of all algorithms considered in this paper is given in Figure 1b when both \mathbf{u} and \mathbf{v} have non-Gaussian priors. Specifically, \mathbf{u} and \mathbf{v} are both distributed IID according to a Rademacher distribution (i.e., each variable is ± 1 with equal probability). We observe that only the AMP algorithm approaches the MMSE (which is invariant to α in this case) and the other methods (gradient descent, joint PCA, and weighted PCA) perform similarly to the Gaussian priors case, which can be seen in Figure 1a. Two additional examples involving non-Gaussian priors are shown in Supplement D.

Gaussian priors in a three group model The final experiment involves $K = 3$ groups with independent standard Gaussian priors. Both equal groups sizes ($\beta_1 = \beta_2 = \beta_3 = 1/3$), and unequal group sizes ($\beta_1 = 0.2, \beta_2 = \beta_3 = 0.4$) are considered. We restrict our attention to Λ matrices such that 1) all the nonzero entries have the same value and 2) the total power satisfies $\beta^\top \Lambda \beta = \lambda$ for some $\lambda > 0$. Under these assumptions, Λ is specified uniquely by the pair (β, λ) and the support (i.e., the indices of the nonzero entries

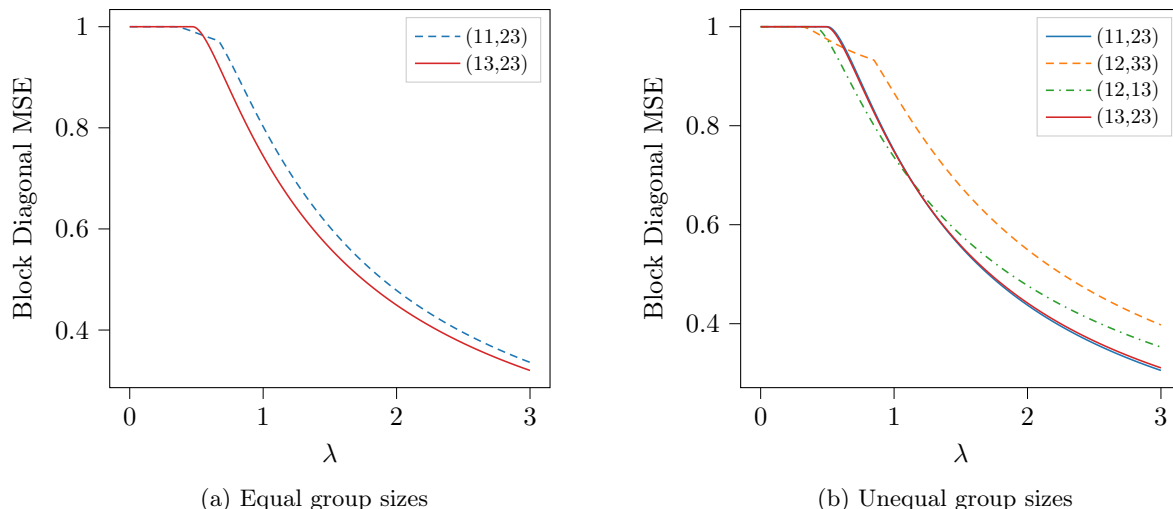


Figure 2: MMSEs for a three group model with Gaussian priors. The legend entry $(ij, k\ell)$ denotes the MMSE for the setup with signal parameters given by a Λ matrix with indices of non-zero entries ij and $k\ell$ and normalized to signal level λ , which is given by the x-axis values on both plots. Note that we only show a single example setup for each unique curve produced, i.e., certain combinations with support size two are omitted because they are redundant.

in Λ). For example, for the support $(11, 23)$ we have

$$\Lambda = \frac{\lambda}{\beta_1^2 + \beta_2\beta_3} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}. \quad (33)$$

Figure 2 compares the asymptotic MMSE, as a function of λ , for all possible supports of size two. Notice that the weak recovery threshold no longer occurs at $\lambda = 1$. Interestingly, the different parameter settings are not ordered uniformly with respect to λ . In other words, the MMSE curves cross each other. One interpretation of these results is that focusing the signal power on interactions without any common factors improves performance at very low signal levels (in fact it lowers the weak recovery threshold) but degrades performance at higher signal levels.

The differences are more pronounced in the case of unequal group sizes (Figure 2b), where the effect of group 1 being substantially smaller than the others dominates the behavior of some of the settings, namely those corresponding to $(11, 23)$ and $(12, 13)$.

5 CONCLUSION

This paper considers a block heterogeneous variation of the extensively studied low-rank matrix estimation problem. The information-theoretic limits for this model are derived by embedding it within the matrix tensor-product model studied by Reeves (2020). We derive single-letter formulas for the MMSE that depend on the noise level in each block and the prior

distributions of the underlying variables. Complementary to the fundamental limits, we also consider the empirical performance of methods based on spectral decomposition, gradient descent, and AMP.

The assumptions used in our theoretical results can be relaxed to a certain extent. For example, following the approach in Reeves (2020), the number of groups K can be allowed to grow at the rate N^α for sufficiently small positive number α and the assumption that the entries are identically distributed can also be relaxed provided that per-variable relative entropy function has a well-defined limit.

An interesting open question for future work is whether a properly defined spectral method can achieve the weak recovery threshold in the two-group model (29) for all values of α .

References

- Abbe, E. (2018). Community detection and stochastic block models: Recent developments. *Journal of Machine Learning Research*, 18(177):1–86.
- Amini, A. A. and Wainwright, M. J. (2009). High-dimensional analysis of semidefinite relaxations for sparse principal components. *Annals of Statistics*, 37(5):2877–2921.
- Arroyo, J., Athreya, A., Cape, J., Chen, G., Priebe, C. E., and Vogelstein, J. T. (2020). Inference for multiple heterogeneous networks with a common invariant subspace.
- Aubin, B., Loureiro, B., Maillard, A., Krzakala, F., and

- Zdeborová, L. (2021). The spiked matrix model with generative priors. *IEEE Transactions on Information Theory*, 67(2):1156–1181.
- Bai, Z. and Yao, J. (2012). On sample eigenvalues in a generalized spiked population model. *Journal of Multivariate Analysis*, 106:167–177.
- Baik, J., Arous, G. B., and Pécché, S. (2005). Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *Annals of Probability*, 33:1643–1697.
- Banks, J., Moore, C., Verzelen, N., Vershynin, R., and Xu, J. (2018). Information-theoretic bounds and phase transitions in clustering, sparse PCA, and submatrix localization. *IEEE Transactions on Information Theory*, 64(7):4872–4894.
- Barbier, J., Dia, M., Macris, N., Krzakala, F., Lesieur, T., and Zdeborová, L. (2016). Mutual information for symmetric rank-one matrix estimation: A proof of the replica formula. In *Advances in Neural Information Processing Systems (NIPS)*, volume 29, pages 424–432, Barcelona, Spain.
- Barbier, J., Dia, M., Macris, N., Krzakala, F., and Zdeborová, L. (2018). Rank-one matrix estimation: analysis of algorithmic and information theoretic limits by the spatial coupling method. *arXiv preprint arXiv:1812.02537*.
- Barbier, J. and Macris, N. (2019). The adaptive interpolation method: a simple scheme to prove replica formulas in bayesian inference. *Probability Theory and Related Fields*, 174(3):1133–185.
- Barbier, J., Macris, N., and Rush, C. (2020). All-or-nothing statistical and computational phase transitions in sparse spiked matrix estimation. In *Advances in Neural Information Processing Systems*, volume 33, pages 14915–14926.
- Barbier, J. and Reeves, G. (2020). Information-theoretic limits of a multiview low-rank symmetric spiked matrix model. In *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*.
- Bayati, M. and Montanari, A. (2011). The dynamics of message passing on dense graphs, with applications to compressed sensing. *IEEE Transactions on Information Theory*, 57(2):764–785.
- Benaych-Georges, F. and Nadakuditi, R. R. (2012). The singular values and vectors of low rank perturbations of large rectangular random matrices. *Journal of Multivariate Analysis*, 111:120–135.
- Bhojanapalli, S., Neyshabur, B., and Srebro, N. (2016). Global optimality of local search for low rank matrix recovery. In *Advances in Neural Information Processing Systems*, volume 29.
- Binkiewicz, N., Vogelstein, J. T., and Rohe, K. (2017). Covariate-assisted spectral clustering. *Biometrika*, 104(2):361–377.
- Chen, H.-B., Mourrat, J.-C., and Xia, J. (2021). Statistical inference of finite-rank tensors. *arXiv preprint arXiv:2104.05360*.
- Cocola, J., Hand, P., and Voroninski, V. (2020). Nonasymptotic guarantees for spiked matrix recovery with generative priors. In *Advances in Neural Information Processing Systems*, volume 33, pages 15185–15197.
- Decelle, A., Krzakala, F., Moore, C., and Zdeborová, L. (2011). Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Physical Review E*, 84(6).
- Deshpande, Y. and Montanari, A. (2014a). Information-theoretically optimal sparse PCA. In *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, pages 2197–2201, Honolulu, HI.
- Deshpande, Y. and Montanari, A. (2014b). Sparse PCA via covariance thresholding. In *Advances in Neural Information Processing Systems*, pages 334–432.
- Deshpande, Y., Montanari, A., Mossel, E., and Sen, S. (2018). Contextual stochastic block models. In *NeurIPS*.
- Fletcher, A. K. and Rangan, S. (2018). Iterative reconstruction of rank-one matrices in noise. *Information and Inference*, 7(3):531–562.
- Guo, D., Shamai, S., and Verdú, S. (2005). Mutual information and minimum mean-square error in Gaussian channels. *IEEE Transactions on Information Theory*, 51(4):1261–1282.
- Guo, D. and Verdú, S. (2005). Randomly spread CDMA: Asymptotics via statistical physics. *IEEE Transactions on Information Theory*, 51(6):1983–2010.
- Guo, D., Wu, Y., Shamai, S., and Verdú, S. (2011). Estimation in Gaussian noise: Properties of the minimum mean-square error. *IEEE Transactions on Information Theory*, 57(4):2371–2385.
- Hiriart-Urruty, J.-B. and Lemaréchal, C. (1993). *Convex Analysis and Minimization Algorithms I*. Springer-Verlag.
- Hong, D., Balzano, L., and Fessler, J. A. (2018a). Asymptotic performance of pca for high-dimensional heteroscedastic data. *Journal of Multivariate Analysis*, 167:435–452.
- Hong, D., Fessler, J. A., and Balzano, L. (2018b). Optimally weighted pca for high-dimensional heteroscedastic data. *arXiv preprint arXiv:1810.12862*.

- Johnstone, I. M. (2001). On the distribution of the largest eigenvalue in principal components analysis. *Annals of Statistics*, 29(2):295–327.
- Johnstone, I. M. and Lu, A. Y. (2009). On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association*, 104(486):682–703.
- Jones, E., Oliphant, T., Peterson, P., et al. (2001–). SciPy: Open source scientific tools for Python.
- Krzakala, F., Xu, J., and Zdeborová, L. (2016). Mutual information in rank-one matrix estimation. In *Proceedings of the IEEE Information Theory Workshop (ITW)*.
- Lelarge, M. and Miolane, L. (2019). Fundamental limits of symmetric low-rank matrix estimation. *Probability Theory and Related Fields*, 173(3):859–929.
- Lesieur, T., de Bacco, C., Banks, J., Krzakala, F., Moore, C., and Zdeborová, L. (2016). Phase transitions and optimal algorithms in high-dimensional gaussian mixture clustering. In *Proceedings of the Allerton Conference on Communication, Control, and Computing*, pages 601–608, Monticello, IL.
- Lesieur, T., Krzakala, F., and Zdeborová, L. (2015). Phase transitions in sparse PCA. In *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, pages 1635–1639, Hong Kong.
- Lesieur, T., Krzakala, F., and Zdeborová, L. (2017). Constrained low-rank matrix estimation: Phase transitions, approximate message passing and applications. *Journal of Statistical Mechanics: Theory and Experiment*, 2017(7).
- Luneau, C., Barbier, J., and Macris, N. (2020). Mutual information for low-rank even-order symmetric tensor estimation. *Information and Inference: A Journal of the IMA*.
- Ma, Z. and Nandy, S. (2021). Community detection with contextual multilayer networkscommunity detection with contextual multilayer networks. *arXiv preprint arXiv:2104.02960*.
- Mayya, V. and Reeves, G. (2019). Mutual information in community detection with covariate information and correlated networks. In *Proceedings of the Allerton Conference on Communication, Control, and Computing*, pages 602–607, Monticello, IL.
- Milgrom, P. and Segal, I. (2002). Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601.
- Miolane, L. (2017). Fundamental limits of low-rank matrix estimation: the non-symmetric case. *arXiv preprint arXiv:1702.00473*.
- Montanari, A. and Venkataramanan, R. (2021). Estimation of low-rank matrices via approximate message passing. *The Annals of Statistics*, 49(1):321 – 345.
- Niles-Weed, J. and Zadik, I. (2020). The all-or-nothing phenomenon in sparse tensor pca. In *Advances in Neural Information Processing Systems*, volume 33, pages 17674–17684.
- Pandit, P., Sahraee Ardakan, M., Rangan, S., Schniter, P., and Fletcher, A. K. (2020). Matrix inference and estimation in multi-layer models. In *Advances in Neural Information Processing Systems*, volume 33, pages 22456–22467.
- Parker, J. T., Schniter, P., and Cevher, V. (2014). Bilinear generalized approximate message passing—part I: Derivation. *IEEE Transactions on Signal Processing*, 62(22):5839–5853.
- Pearl, J. (1998). *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- Perry, A., Wein, A. S., Bandeira, A. S., and Moitra, A. (2018). Optimality and sub-optimality of PCA I: Spiked random matrix models. *Annals of Statistics*, 46(5):2416–2451.
- Polyanskiy, Y. and Wu, Y. (2016). Wasserstein continuity of entropy and outer bounds for interference channels. *IEEE Transactions on Information Theory*, 62(7):3992–4002.
- Reeves, G. (2020). Information-theoretic limits for the matrix tensor product. *IEEE Journal on Selected Areas in Information Theory*, 1(3):777–798.
- Rockafellar, R. T. (1970). *Convex Analysis*. Princeton University Press.
- Yao, J., Zheng, S., and Bai, Z. (2015). *Large Sample Covariance Matrices and High-Dimensional Data Analysis*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.
- Zdeborová, L. and Krzakala, F. (2016). Statistical physics of inference: Thresholds and algorithms. *Advances in Physics*, 65(5):453–552.
- Zhang, A. R., Cai, T. T., and Wu, Y. (2021). Heteroskedastic pca: Algorithm, optimality, and applications. *arXiv preprint arXiv:1810.08316*.
- Zhang, J. and Zhang, R. (2020). How many samples is a good initial point worth in low-rank matrix recovery? In *Advances in Neural Information Processing Systems*, volume 33, pages 12583–12592.
- Zou, H., Hastie, T., and Tibshirani, R. (2006). Sparse principal component analysis. *Journal of Computational and Graphical Statistics*, 15(2):265–286.

Fundamental limits for rank-one matrix estimation with groupwise heteroskedasticity: Supplementary Materials

A PROOFS OF MAIN RESULTS

As described in Section 2.1 the main idea in the proof is to show that the groupwise model in (1) can be expressed as a special case of the matrix tensor product model studied in Reeves (2020). In the following, we first summarize the basic definition and results from Reeves (2020) and then provide the proofs for Theorems 1 and 2.

A.1 The Matrix Tensor Product Model

Let X be an $n \times d$ random matrix. Following (Reeves, 2020, Equation (9)), the matrix tensor product model associated with a positive scale factor N , $d \times d$, symmetric positive semidefinite matrix R , and $d^2 \times d^2$ symmetric positive semidefinite matrix S is defined by

$$Y_{R,S}^{\text{MTP}} = \begin{cases} XR^{1/2} + W \\ \frac{1}{\sqrt{N}}(X \otimes X)S^{1/2} + W' \end{cases} \quad (34)$$

where \otimes denotes the Kronecker matrix product and W and W' are independent matrices with independent standard Gaussian entries. The only difference between this definition and the one in Reeves (2020) is that the scale factor N maybe different from the number of rows n . The relative entropy function associated with this model is defined by

$$\mathcal{D}_N^{\text{MTP}}(R, S) := \frac{1}{N} D(\mathbb{P}_{Y_{R,S}^{\text{MTP}}} \| \mathbb{P}_{Y_{0,0}^{\text{MTP}}}), \quad (35)$$

where $\mathbb{P}_{Y_{0,0}^{\text{MTP}}}$ is the standard Gaussian measure of the same dimension as the observations. The approximation formula (see (Reeves, 2020, Equation (42))) is defined as

$$\hat{\mathcal{D}}_N^{\text{MTP}}(R, S) := \max_Q \inf_{\tilde{R} \in \mathcal{S}_+^d} \left\{ \mathcal{D}_N^{\text{MTP}}(R + \tilde{R}, 0) + \frac{1}{2} \text{tr}(S(Q \otimes Q)) - \frac{1}{2} \text{tr}(\tilde{R}Q) \right\}. \quad (36)$$

where the maximum is over $d \times d$ symmetric positive semidefinite matrices satisfying $Q \preceq \frac{1}{N} \mathbb{E}[X^\top X]$.

The main results in Reeves (2020) are bounds on the difference between the relative entropy and the approximation formula. For the purposes of this paper, we are interested in the asymptotic setting where the number of rows increases while the number of columns is fixed.

Theorem 5 (Reeves (2020)). *Consider a sequence of problems, indexed by integer N , with observations from the matrix tensor product model (34) such that the number of columns d is fixed while the number of rows scales as $n/N \rightarrow \beta \in (0, \infty)$. If the rows of X are independent with $\mathbb{E}[|X_{ij}|^4] \leq C$ for some constant C , then, $|\mathcal{D}_N^{\text{MTP}} - \hat{\mathcal{D}}_N^{\text{MTP}}|$ converges to zero pointwise as $N \rightarrow \infty$.*

Proof. This result follows along similar lines as (Reeves, 2020, Corollary 1). If the entries of X are bounded uniformly, then the result follows directly from (Reeves, 2020, Theorem 1). The relaxation from bounded entries to bounded fourth moment follows from the continuity of relative entropy with respect to the quadratic Wasserstein distance Polyanskiy and Wu (2016) and the fact that a finite fourth moment bound on the entries of X implies a finite second moment bound on the entries of $X \otimes X$. \square

A.2 Proof of Theorem 1

Let $X = \text{diag}(\mathbf{x}_1, \dots, \mathbf{x}_K)$ be the $n \times K$ block-diagonal matrix whose k -th block is given by \mathbf{x}_k and let Γ be the $K \times K$ matrix corresponding to the entrywise positive square root of Λ . Using this notation, the groupwise model (1) can be expressed in matrix form according to

$$Y = \frac{1}{\sqrt{N}} X \Gamma X^\top + W. \quad (37)$$

By vectorization, this model can be written equivalently as

$$\text{vec}(Y) = \frac{1}{\sqrt{N}} (X \otimes X) \text{vec}(\Gamma) + \text{vec}(W). \quad (38)$$

This is the matrix tensor product model given in (Reeves, 2020, Definition 1) with $K^2 \times 1$ coupling matrix $\text{vec}(\Gamma)$. As noted in (Reeves, 2020, Proposition 1), the fundamental limits for this model are the same as for the symmetric form of the matrix tensor product model (see (Reeves, 2020, Definition 2)) characterized by $K^2 \times K^2$ coupling matrix $\text{vec}(\Gamma) \text{vec}(\Gamma)$. Recalling that the relative entropy function $\mathcal{D}_N(r, \Lambda)$ in (8) is defined with respect to observations from the groupwise model (1) and the side-information model (6) we can write

$$\mathcal{D}_N(r, \Lambda) = \mathcal{D}_N^{\text{MTP}}(\text{diag}(r), \text{vec}(\Gamma) \text{vec}(\Gamma)^\top) \quad (39)$$

where $\mathcal{D}_N^{\text{MTP}}(R, S)$ is defined in (35).

From here, a direct application of Theorem 5 shows that the limiting behavior of \mathcal{D}_N can be expressed in terms of the approximation formula for the matrix tensor product in (36). To simplify the analysis, we use the following result, which leverages the block-diagonal structure in X .

Lemma 6. *Suppose that $X = \text{diag}(\mathbf{x}_1, \dots, \mathbf{x}_K)$ is a block-diagonal matrix where each \mathbf{x}_k is a column vector. Then, the relative entropy function in the matrix tensor product model depends only on the diagonal entries of the matrices (R, S) . Specifically,*

$$\mathcal{D}_N^{\text{MTP}}(R, S) = \mathcal{D}_N^{\text{MTP}}(\text{diag}(r), \text{diag}(s)), \quad (40)$$

where r and s are vectors containing the diagonal entries in R and S , respectively.

Proof. Consider the observation model $A\mathbf{u} + \mathbf{w}$ where \mathbf{u} is a $p \times 1$ random vector, $A = [a_1, \dots, a_K]$ is an $m \times p$ deterministic matrix, and \mathbf{w} is an independent $m \times 1$ vector with standard Gaussian entries. Suppose that the entries of \mathbf{u} are deterministically zero except in the k -th location and let Q_k be an $m \times m$ orthogonal matrix (chosen independently of \mathbf{u} and \mathbf{w}) such that the k -th row of Q_k is given by $a_k / \|a_k\|$. Then, we can write

$$D(A\mathbf{u} + \mathbf{w} \parallel \mathbf{w}) = D(a_k u_k + \mathbf{w} \parallel \mathbf{w}) \quad (41)$$

$$= D(Q_k a_k u_k + Q_k \mathbf{w} \parallel Q_k \mathbf{w}) \quad (42)$$

$$= D((\text{diag}(\|a_1\|, \dots, \|a_K\|)\mathbf{u} + Q_k \mathbf{w} \parallel Q_k \mathbf{w}) \quad (43)$$

where the second line follows from the invariance of relative entropy to one-to-one transformation and the third line follows from first noting that $Q_k a_k = \|a_k\| e_k$, where e_k is the k -th standard basis vector, and then recalling that only the k th entry in \mathbf{u} is nonzero. By the orthogonal invariance of the standard Gaussian distribution $Q_k \mathbf{w}$ is a standard Gaussian vector that is independent of \mathbf{u} . Now suppose that T is a $p \times p$ symmetric positive semidefinite matrix with diagonal given by t . Setting $A = T^{1/2}$ we find that $(\|a_1\|^2, \dots, \|a_K\|^2) = t$ and so the arguments given above imply that

$$D(T^{1/2}\mathbf{u} + \mathbf{w} \parallel \mathbf{w}) = D(\text{diag}(t)^{1/2}\mathbf{u} + \mathbf{w} \parallel \mathbf{w}). \quad (44)$$

By the chain rule for relative entropy, this result extends to the setting where U is an $n \times p$ random matrix with the property that each row has at most one entry that is not deterministically zero, i.e., we can write

$$D(UT^{1/2} + W \parallel W) = D(U \text{diag}(t)^{1/2} + W \parallel W). \quad (45)$$

where W is an independent $n \times p$ matrix with independent standard Gaussian entries. Finally, the stated result for the matrix tensor product model follows from noting that both X and $X \otimes X$ have the property that at most one entry in each row is not deterministically zero. \square

In view of Lemma 6, the relation in (39) can be simplified as

$$\mathcal{D}_N(r, \Lambda) = \mathcal{D}_N^{\text{MTP}}(\text{diag}(r), \text{diag}(\text{vec}(\Lambda))). \quad (46)$$

Moreover, if \tilde{R} is a $K \times K$ symmetric positive semidefinite matrix with diagonal given by \tilde{r} , we have

$$\mathcal{D}_N^{\text{MTP}}(\text{diag}(r) + \tilde{R}, 0) = \mathcal{D}_N^{\text{MTP}}(\text{diag}(r_k + \tilde{r}), 0) = \sum_{k=1}^K \frac{n_k}{N} D_k(r_k + \tilde{r}_k) \quad (47)$$

where the first step is due to Lemma 6 and second step follows from Assumption 1. Combining this expression with the fact that $\frac{1}{N} \mathbb{E}[X^\top X] = \text{diag}(n_1/N, \dots, n_K/N)$, leads to the following simplification of approximation formula:

$$\hat{\mathcal{D}}_N^{\text{MTP}}(\text{diag}(r), \text{diag}(\text{vec}(\Lambda))) = \max_q \inf_{\tilde{r}} \left\{ \sum_{k=1}^K \frac{n_k}{N} D_k(r_k + \tilde{r}_k) + \frac{1}{2} q^\top \Lambda q - \frac{1}{2} \tilde{r}^\top q \right\}. \quad (48)$$

Here, the maximum is over K -dimensional vectors satisfying $0 \leq q_k \leq n_k/N$ and the infimum is over K -dimensional vectors with non-negative entries. Finally, taking the $N \rightarrow \infty$ limit of this expression leads to the formula for the limit given in Theorem 1.

A.3 Proof of Theorem 2

The proof of Theorem 2 follows from fact that changes in the relative entropy with respect to the parameters (r, Λ) can be related to the MMSE in estimating the underlying variables. As a consequence of the I-MMSE relationship Guo and Verdú (2005), it can be shown that \mathcal{D}_N is differentiable on the interior of its domain with partial derivatives:

$$\partial_{r_k} \mathcal{D}_N(r, \Lambda) = \frac{1}{2N} \mathbb{E}[\|\mathbb{E}[\mathbf{x}_k | Y]\|_2^2], \quad \partial_{\lambda_{k\ell}} \mathcal{D}_N(r, \Lambda) = \frac{1}{2N^2} \mathbb{E}[\|\mathbb{E}[\mathbf{x}_k \mathbf{x}_\ell^\top | Y]\|_F^2]. \quad (49)$$

These derivatives can be defined uniquely on the boundaries ($r_k = 0$ or $\lambda_{k\ell} = 0$) according to the right derivatives, which exist and are finite due the finite fourth moment assumption (Guo et al., 2011, Proposition 7). The connection with the MMSE follows from the relationship

$$\mathbb{E}[\|\mathbf{x}_k - \mathbb{E}[\mathbf{x}_k | Y]\|^2] = \mathbb{E}[\|\mathbf{x}_k\|^2] - \mathbb{E}[\|\mathbb{E}[\mathbf{x}_k | Y]\|_2^2] \quad (50)$$

$$\mathbb{E}[\|\mathbf{x}_k \mathbf{x}_\ell^\top - \mathbb{E}[\mathbf{x}_k \mathbf{x}_\ell^\top | Y]\|_F^2] = \mathbb{E}[\|\mathbf{x}_k \mathbf{x}_\ell^\top\|_F^2] - \mathbb{E}[\|\mathbb{E}[\mathbf{x}_k \mathbf{x}_\ell^\top | Y]\|_F^2]. \quad (51)$$

The assumption that the entries are independent with second moment equal to one means that $\mathbb{E}[\|\mathbf{x}_k\|^2] = n_k$ and $\mathbb{E}[\|\mathbf{x}_k \mathbf{x}_\ell^\top\|_F^2] = n_k n_\ell$ for $k \neq \ell$. For the diagonal blocks, we can write $\mathbb{E}[\|\mathbf{x}_k \mathbf{x}_k^\top\|_F^2] = n_k \xi_k + n_k(n_k - 1)$ where ξ_k is the fourth moment of P_k and so $n_k^{-2} \mathbb{E}[\|\mathbf{x}_k \mathbf{x}_k^\top\|_F^2]$ converges to one in the $N \rightarrow \infty$ limit.

The next step is to relate the derivatives of \mathcal{D}_N to the derivatives of its limit \mathcal{D} . This is possible due to the fact that \mathcal{D}_N is convex, which implies that \mathcal{D} is convex, and hence differentiable almost everywhere, and furthermore that the derivatives converge at every point where the limit is differentiable (Hiriart-Urruty and Lemaréchal, 1993, Proposition 4.3.4).

With these results in hand, the asymptotic MMSE can be analyzed by studying the limit \mathcal{D} . If for a given pair (r, Λ) the maximum in the definition of \mathcal{D} is attained at a unique point q^* then \mathcal{D} has partial derivatives with respect to the nonzero entries in (r, Λ) . Specifically, for $r_k > 0$ and $\lambda_{k\ell} > 0$ it can be verified that

$$\partial_{r_k} \mathcal{D}(r, \Lambda) = \frac{1}{2} q_k^*, \quad \partial_{\lambda_{k\ell}} \mathcal{D}(r, \Lambda) = \frac{1}{2} q_k^* q_\ell^*. \quad (52)$$

For $r = 0$ and $\lambda_{k\ell} = 0$ these expressions still hold for the right derivatives and it can be shown (see (Reeves, 2020, Section II-B)) that they provide one-sided bounds of the form

$$\limsup_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}[\|\mathbb{E}[\mathbf{x}_k | Y]\|_2^2] \leq q_k^*, \quad \limsup_{N \rightarrow \infty} \frac{1}{N^2} \mathbb{E}[\|\mathbb{E}[\mathbf{x}_k \mathbf{x}_\ell^\top | Y]\|_F^2] \leq q_k^* q_\ell^*. \quad (53)$$

Noting that these are equalities whenever the upper bound is equal to zero and combining with the arguments above leads to the conditions in Theorem 2.

A.4 Almost everywhere uniqueness of maximizer

Recall that \mathcal{D} is convex and thus differentiable almost everywhere on the interior of its domain (Rockafellar, 1970, Theorem 25.5). For each point (r, Λ) where \mathcal{D} is differentiable, we can combine the envelop theorem Milgrom and Segal (2002) with (52) to conclude that the maximum in (9) has a unique maximizer q^* . Finally, this argument can also be applied with respect to a subset of entries in (r, Λ) , providing that the corresponding partial derivatives uniquely characterize q^* . For example, for any fixed r and collection of off-diagonal terms $(\lambda_{k\ell} : k \neq \ell)$, the maximizer is unique for almost all $\lambda_{11}, \dots, \lambda_{KK} \in (0, \infty)$.

A.5 Proof of Theorem 4

Recall that the relative entropy function given in Theorem 1 for the optimally weighted PCA setting described in Section 3 is given by

$$\mathcal{D}(0, \Lambda) = \max_{0 \leq q \leq \beta} \inf_{\tilde{r} \geq 0} \left\{ \sum_{k=0}^L \beta_k D_k(\tilde{r}_k) + \frac{q_0}{2} \sum_{\ell=1}^L \frac{q_\ell}{\sigma_\ell^2} - \frac{1}{2} \tilde{r}^\top q \right\} \quad (54)$$

where $D_k : [0, \infty) \rightarrow [0, \infty)$ is the single letter relative entropy function for $k = 0, \dots, L$. The assumption that the entries of \mathbf{u} and \mathbf{v}_ℓ have independent standard Gaussian entries means that the single-letter relative entropy functions can be expressed in closed-form as $D_k(\tilde{r}_k) = \frac{1}{2}(\tilde{r}_k - \log(1 + \tilde{r}_k))$ for $k = 0, 1, \dots, L$. Observe that the objective function in (54) is convex in \tilde{r} . By differentiation, one finds that for $0 \leq q < \beta$, the infimum is attained at a unique point $r^* = r^*(q)$ satisfying:

$$r_k^* = \frac{q_k}{\beta_k - q_k}, \quad k = 0, 1, \dots, L \quad (55)$$

Plugging this value into (54) and noting that the objective is continuous in q leads to the following expression:

$$\sup_{0 \leq q < \beta} \left\{ \frac{q_0}{2} \left(1 + \sum_{\ell=1}^L \frac{q_\ell}{\sigma_\ell^2} \right) + \frac{1}{2} \sum_{\ell=1}^L q_\ell - \frac{1}{2} \sum_{k=0}^L \beta_k \log(\beta_k) + \frac{1}{2} \sum_{k=0}^L \beta_k \log(\beta_k - q_k) \right\} \quad (56)$$

This expression can be simplified further by noting that for fixed $0 \leq q_0 < \beta_0$ the objective function in (56) is strictly concave in (q_1, \dots, q_L) and thus the maximum is attained at the unique point $(q_1^*(q_0), \dots, q_L^*(q_0))$ characterized by

$$q_\ell^* = \frac{\beta_\ell q_0}{\sigma_\ell^2 + q_0}, \quad \ell = 1, \dots, L. \quad (57)$$

In the case $q_0 = 0$ this follows since the objective function is decreasing in q_ℓ for $\ell \geq 1$ and thus the maximum is attained at the boundary point $q_1^* = \dots = q_L^* = 0$. In the case $q_0 > 0$ this follows from differentiation. Evaluating (56) at the point given by (57) and then simplifying terms leads to the following expression involving a univariate optimization problem:

$$\sup_{0 \leq q_0 < \beta_0} \frac{1}{2} F(q_0), \quad F(x) := x - \beta_0 \log\left(\frac{\beta_0}{\beta_0 - x}\right) + \sum_{\ell=1}^L \left(\frac{\beta_\ell x}{\sigma_\ell^2} - \beta_\ell \log\left(1 + \frac{x}{\sigma_\ell^2}\right) \right) \quad (58)$$

The objective function F is continuous on $[0, \beta_0)$ and converges to $-\infty$ at $x \nearrow \beta_0$. Hence, the maximum is attained on $[0, \beta_0)$. Furthermore, the derivative can be expressed as

$$F'(x) = -\frac{x}{\beta_0 - x} + \sum_{\ell=1}^L \frac{\beta_\ell x}{\sigma_\ell^2(\sigma_\ell^2 + x)}. \quad (59)$$

From inspection, it is clear that F' has a root at $x = 0$. To explore the possible roots on $(0, \beta_0)$ define

$$R(x) := \left(1 - \frac{\beta_0}{x}\right) F'(x) = 1 - \sum_{\ell=1}^L \frac{\beta_\ell(\beta_0 - x)}{\sigma_\ell^2(\sigma_\ell^2 + x)} \quad (60)$$

Note that $(1 - \beta_0/x)$ is strictly negative on $(0, \beta_0)$ and so R and F' have the same roots and opposite signs over this interval.

At this point we recognize that R is the same as function appearing in Hong et al. (2018b, Theorem 6), adapted to the notation used in this paper. The analysis in Hong et al. (2018b) shows that R is strictly increasing over $(0, \beta_0)$ and converges to 1 as $x \nearrow \beta_0$. Consequently, it follows from (60) that there are only two cases to consider:

- If $R(0) \geq 0$ then F' is strictly negative on $(0, \beta_0)$, and so the unique maximizer is $q_0^* = 0$.
- Conversely, if $R(0) < 0$, then there exists a single root x^* on $(0, \beta_0)$. Moreover, F' is strictly positive on $(0, x^*)$ and strictly negative on (x^*, β_0) and thus $q_0^* = x^*$ is the unique maximizer of F .

To conclude, observe the condition on $R(0)$ can be expressed equivalently as

$$R(0) < 0 \iff \sum_{\ell=1}^L \frac{\beta_0 \beta_\ell}{\sigma_\ell^4} > 1 \tag{61}$$

The above establishes Theorem 4 and now we can use it in conjunction with Theorem 2 and Theorem 3 to comment on the optimality of WPCA in this setting. First recall that Hong et al. (2018b, Theorem 6) states that

$$\frac{\langle \hat{\mathbf{u}}, \mathbf{u} \rangle^2}{\|\mathbf{u}\|^2 \|\hat{\mathbf{u}}\|^2} \xrightarrow{a.s.} q_0^{\text{WPCA}} / \beta_0 \tag{62}$$

where $\mathbf{u}, \hat{\mathbf{u}} \in \mathbb{R}^d$ and $\hat{\mathbf{u}}$ is the estimate produced by WPCA with optimal weights. It should be noted that Hong et al. (2018b) assume that \mathbf{u} is a unit vector, however we take it to be a vector with IID entries with zero mean and unit variance. This discrepancy leads to some additional normalization terms in the formulas presented here.

We consider the MSE in estimating $\mathbf{u}\mathbf{u}^\top$ and allow for the optimal scaling of $\hat{\mathbf{u}}$. We recall that the formula for the MSE in this case was derived in Section 4 and is

$$\inf_{\rho_u \geq 0} \frac{1}{d^2} \mathbb{E} \left[\left\| \mathbf{u}\mathbf{u}^\top - \rho_u^2 \frac{\hat{\mathbf{u}}\hat{\mathbf{u}}^\top}{\|\hat{\mathbf{u}}\|^2} \right\|_F^2 \right] = \frac{1}{d^2} \mathbb{E}[\|\mathbf{u}\|^4] - \frac{1}{d^2} \mathbb{E} \left[\frac{\langle \mathbf{u}, \hat{\mathbf{u}} \rangle^2}{\|\hat{\mathbf{u}}\|^2} \right]^2. \tag{63}$$

We can consider the above MSE in the asymptotic limit (i.e., as $N \rightarrow \infty$) and note that in this limit $\|\mathbf{u}\|^2$ concentrates to a deterministic value of d . So, the asymptotic MSE for WPCA is given by

$$\lim_{N \rightarrow \infty} \frac{1}{d^2} \left(\mathbb{E}[\|\mathbf{u}\|^4] - \left(\mathbb{E} \left[\frac{\langle \mathbf{u}, \hat{\mathbf{u}} \rangle^2}{\|\hat{\mathbf{u}}\|^2} \right] \right)^2 \right) = 1 - \lim_{N \rightarrow \infty} \left(\mathbb{E} \left[\frac{\langle \mathbf{u}, \hat{\mathbf{u}} \rangle^2}{\|\mathbf{u}\|^2 \|\hat{\mathbf{u}}\|^2} \right] \right)^2 = 1 - \left(\frac{q_0^{\text{WPCA}}}{\beta_0} \right)^2 \tag{64}$$

where the final equality holds because

$$\frac{\langle \mathbf{u}, \hat{\mathbf{u}} \rangle^2}{\|\mathbf{u}\|^2 \|\hat{\mathbf{u}}\|^2} \tag{65}$$

is a positive and bounded quantity. Therefore, almost sure convergence implies convergence in L_1 . We also note that this MSE is lower bounded by the MMSE. Furthermore, from Theorem 3, the MMSE of $\mathbf{u}\mathbf{u}^\top$ is lower bounded by

$$1 - \left(\frac{q_0^*}{\beta_0} \right)^2 \tag{66}$$

where q_0^* is the result given in Theorem 4 and is equal q_0^{WPCA} at and after the weak recovery condition is satisfied. Given that we have demonstrated an algorithm that attains a lower bound on the MMSE, we can conclude that the MSE of the optimally scaled WPCA estimator is equal to the MMSE. Hence, it can be said that WPCA achieves the fundamental limits in this setting.

B AMP FOR HETEROSKEDASTIC RANK-ONE ESTIMATION

This section considers approximate inference methods for the heteroskedastic rank-one observation model given by

$$y_{ij} = \sqrt{\frac{\lambda_{ij}}{N}} x_i x_j + w_{ij}, \quad i, j = 1, \dots, n \quad (67)$$

where (x_i) are unknown random variables, (w_{ij}) are independent standard Gaussian noise terms, and (λ_{ij}) and N are known parameters. Furthermore, it is assumed that the variables (x_i) are independently (but not identically) distributed with $x_i \sim P_i$ for $i = 1, \dots, n$. Note that the groupwise model (1) corresponds to the special case where the parameters (λ_{ij}) and the distributions are constant within groups.

Symmetrization and Preprocessing To simplify the exposition of the algorithms it is convenient to apply a preprocessing step that re-scales and symmetrizes the observations. Specifically, we define the modified observations (\tilde{y}_{ij}) and parameters $(\tilde{\lambda}_{ij})$ according to

$$\tilde{y}_{ij} = \begin{cases} \sqrt{\lambda_{ij}} y_{ij} + \sqrt{\lambda_{ji}} y_{ji} & i \neq j \\ \sqrt{\lambda_{ii}} y_{ii}, & i = j \end{cases}, \quad \tilde{\lambda}_{ij} = \begin{cases} \lambda_{ij} + \lambda_{ji}, & i \neq j \\ \lambda_{ii} & i = j \end{cases} \quad (68)$$

By the orthogonal invariance of the standard Gaussian distribution it follows that (\tilde{y}_{ij}) are sufficient statistics for estimation of (x_i) . Furthermore, the distribution of (\tilde{y}_{ij}) is given by the symmetric model

$$\tilde{y}_{ij} = \frac{1}{\sqrt{N}} \tilde{\lambda}_{ij} x_i x_j + \tilde{\lambda}_{ij}^{1/2} \tilde{w}_{ij} \quad (69)$$

where $\tilde{w}_{ij} = \tilde{w}_{ji}$ with $(\tilde{w}_{ij} : 1 \leq i < j \leq n)$ independent standard Gaussian variables. Using this parameterization, the likelihood function is proportional to

$$\prod_{1 \leq i < j \leq n} \exp \left\{ \frac{\tilde{y}_{ij} x_i x_j}{\sqrt{N}} - \frac{\tilde{\lambda}_{ij} x_i^2 x_j^2}{2N} \right\}. \quad (70)$$

Belief Propagation We begin with a derivation of the belief propagation algorithm Pearl (1998). To simplify the derivation we assume that each probability measure P_i had a density f_i with respect to the Lebesgue measure. In view of (70), the conditional distribution of the variables (x_i) given the observations (\tilde{y}_{ij}) has a density that is proportional to

$$\prod_{i=1}^n \phi_i(x_i) \prod_{1 \leq i < j \leq n} \psi_{ij}(x_i, x_j), \quad (71)$$

where

$$\phi_i(x_i) := f_i(x_i) \exp \left\{ \frac{\tilde{y}_{ii} x_i^2}{\sqrt{N}} - \frac{\tilde{\lambda}_{ii} x_i^4}{2N} \right\}, \quad \psi_{ij}(x_i, x_j) := \exp \left\{ \frac{\tilde{y}_{ij} x_i x_j}{\sqrt{N}} - \frac{\tilde{\lambda}_{ij} x_i^2 x_j^2}{2N} \right\} \quad (72)$$

The belief propagation algorithm is described by a collection of messages $(\mu_{ij} : 1 \leq i, j \leq n, i \neq j)$ where each μ_{ij} is a non-negative function that describes the ‘influence’ of variable i on variable j . Starting at a given initialization, these messages are determined recursively via the update rule

$$\mu_{ij}(x_j) \leftarrow \int \phi_i(x_i) \psi_{ij}(x_i, x_j) \prod_{k \in [n] \setminus \{i, j\}} \mu_{ki}(x_i) dx_i, \quad (73)$$

where $[n] = \{1, \dots, n\}$. The approximation to the marginal of the i -th variable associated with given collection of messages is given by the probability density function:

$$x_i \mapsto \frac{\phi_i(x_i) \prod_{j \in [n] \setminus \{i\}} \mu_{ji}(x_i)}{\int \phi_i(x'_i) \prod_{j \in [n] \setminus \{i\}} \mu_{ji}(x'_i) dx'_i}. \quad (74)$$

Note that this final approximation is invariant to rescaling of the messages, and thus at each stage in the algorithm the message μ_{ij} can be rescaled by an arbitrary positive constant.

Relaxed Belief Propagation Next we consider a relaxed belief propagation algorithm that is obtained by an approximation of the update rule. The first step in this approximation is to replace the function $\psi_{ij}(x_i, x_j)$ by its second order expansion with respect to the term $(x_i x_j)/\sqrt{N}$ evaluated at zero:

$$\psi_{ij}(x_i, x_j) \approx 1 + \frac{\tilde{y}_{ij} x_i x_j}{\sqrt{N}} + \frac{1}{2} \frac{(\tilde{y}_{ij}^2 - \tilde{\lambda}_{ij}) x_i^2 x_j^2}{N}. \quad (75)$$

Plugging this approximation into (73) leads to the following integral

$$\int \phi_i(x_i) \left(1 + \frac{\tilde{y}_{ij} x_i x_j}{\sqrt{N}} + \frac{1}{2} \frac{(\tilde{y}_{ij}^2 - \tilde{\lambda}_{ij}) x_i^2 x_j^2}{N} \right) \prod_{k \in [n] \setminus \{i, j\}} \mu_{ki}(x_i) dx_i \quad (76)$$

To characterize the solution to this integral, define

$$m_{ij} = \frac{\int x_i \phi_i(x) \prod_{k \in [n] \setminus \{i, j\}} \mu_{ki}(x_i) dx_i}{\int \phi_i(x') \prod_{k \in [n] \setminus \{i, j\}} \mu_{ki}(x'_i) dx'_i} \quad (77)$$

$$v_{ij} = \frac{\int x_i^2 \phi_i(x) \prod_{k \in [n] \setminus \{i, j\}} \mu_{ki}(x_i) dx_i}{\int \phi_i(x') \prod_{k \in [n] \setminus \{i, j\}} \mu_{k \rightarrow i}(x'_i) dx'_i} - \hat{x}_{ij}^2 \quad (78)$$

to be the mean and variance, respectively, of the probability density function

$$x_i \mapsto \frac{\phi_i(x) \prod_{k \in [n] \setminus \{i, j\}} \mu_{ki}(x_i)}{\int \phi_i(x') \prod_{k \in [n] \setminus \{i, j\}} \mu_{ki}(x'_i) dx'_i}. \quad (79)$$

This density is similar to the marginal approximation in (74) except that message from the j -th variable is also excluded. Using this notation, the integral in (76) is proportional to

$$1 + \left(\frac{\tilde{y}_{ij} m_{ij}}{\sqrt{N}} \right) x_j + \frac{1}{2} \left(\frac{(\tilde{y}_{ij}^2 - \tilde{\lambda}_{ij})(m_{ij}^2 + v_{ij})}{N} \right) x_j^2. \quad (80)$$

Using the approximation $1 + au + \frac{1}{2}(a^2 - b)u^2 \approx \exp(au - \frac{1}{2}bu^2)$ for $u \approx 0$ and recalling that we can rescale the message μ_{ij} by an arbitrary positive constant leads to the modified update rule

$$\mu_{ij}(x_j) \leftarrow \exp \left\{ a_{ij} x_j - \frac{1}{2} b_{ij} x_j^2 \right\} \quad (81)$$

where

$$a_{ij} = \frac{\tilde{y}_{ij} m_{ij}}{\sqrt{N}}, \quad b_{ij} = \frac{\tilde{\lambda}_{ij} m_{ij}^2 + (\tilde{\lambda}_{ij} - \tilde{y}_{ij}^2) v_{ij}}{N}. \quad (82)$$

According to the update rule (81), the probability density in (79) is defined by the (unnormalized) density

$$\phi_i(x_i) \prod_{k \in [n] \setminus \{i, j\}} \exp \left\{ a_{ki} x_i - \frac{1}{2} b_{ki} x_i^2 \right\} = \phi_i(x_i) \exp \left\{ \sum_{k \in [n] \setminus \{i, j\}} a_{ki} x_i - \frac{1}{2} b_{ki} x_i^2 \right\} \quad (83)$$

Here, we see that the products of the messages are described compactly in terms of a linear combination of the a_{ki} and b_{ki} terms.

The relaxed belief propagation can be described as follows. For $i = 1, \dots, n$ define the function

$$\eta_i(a, b) = \frac{\int x_i \exp \left\{ a x_i - \frac{1}{2} b x_i^2 \right\} \phi_i(x_i) dx_i}{\int \exp \left\{ a x'_i - \frac{1}{2} b (x'_i)^2 \right\} \phi_i(x'_i) dx'_i} \quad (84)$$

and let $\eta'_i(a, b)$ denote the partial derivative with respect to the first argument. Note that $\eta_i(a, b)$ and $\eta'_i(a, b)$ represent the mean and variance, respectively, of the (unnormalized) density $x_i \mapsto \phi_i(x_i) \exp(ax_i - \frac{1}{2}bx_i^2)$. Starting with initial values for (m_{ij}, v_{ij}) the updates are defined according to

$$m_{ij} \leftarrow \eta_i \left(\sum_{k \in [n] \setminus \{i, j\}} a_{ki}, \sum_{k \in [n] \setminus \{i, j\}} b_{ki} \right) \quad (85a)$$

$$v_{ij} \leftarrow \eta'_i \left(\sum_{k \in [n] \setminus \{i, j\}} a_{ki}, \sum_{k \in [n] \setminus \{i, j\}} b_{ki} \right) \quad (85b)$$

where the terms (a_{ki}, b_{ki}) are defined with respect to the current values of (m_{ki}, v_{ki}) according to (82).

Approximate Message Passing One of the limitations of the relaxed belief propagation algorithm in (85) is that it requires keeping track of n^2 terms. AMP can be viewed as an approximation to the relaxed belief propagation algorithm that requires only n terms.

Before stating the AMP algorithm, it is useful to consider the version of relaxed belief propagation where all of the updates are made in parallel. Specially, starting with an initialization (m_{ij}^0, v_{ij}^0) , the updates at time $t = 1, 2, \dots$ are given by

$$a_{ij}^t = \frac{\tilde{y}_{ij}}{\sqrt{N}} m_{ij}^t \quad (86a)$$

$$b_{ij}^t = \frac{\tilde{\lambda}_{ij}}{N} (m_{ij}^t)^2 + \frac{\tilde{\lambda}_{ki} - \tilde{y}_{ij}}{N} v_{ij}^t \quad (86b)$$

$$m_{ij}^{t+1} = \eta_i \left(\sum_{k \in [n] \setminus \{i, j\}} a_{ki}^t, \sum_{k \in [n] \setminus \{i, j\}} b_{ki}^t \right) \quad (86c)$$

$$v_{ij}^{t+1} = \eta'_i \left(\sum_{k \in [n] \setminus \{i, j\}} a_{ki}^t, \sum_{k \in [n] \setminus \{i, j\}} b_{ki}^t \right) \quad (86d)$$

In comparison, the corresponding version of AMP is given by

$$a_i^t = \frac{1}{\sqrt{N}} \sum_{k \in [n] \setminus i} \tilde{y}_{ki} m_k^t - \frac{1}{N} m_i^{t-1} \sum_{k \in [n] \setminus i} \tilde{y}_{ki} v_k^t \tilde{y}_{ik} \quad (87a)$$

$$b_i^t = \frac{1}{N} \sum_{k \in [n] \setminus i} \left(\tilde{\lambda}_{ki} (m_k^t)^2 + (\tilde{\lambda}_{ki} - \tilde{y}_{ki}^2) v_k^t \right) \quad (87b)$$

$$m_i^{t+1} = \eta_i(a_i^t, b_i^t) \quad (87c)$$

$$v_i^{t+1} = \eta'_i(a_i^t, b_i^t) \quad (87d)$$

The relationship between these algorithms can be derived heuristically based on a decomposition of the means of the form $m_{ij}^t = m_i^t + \epsilon_{ij}^t$ where m_i^t does not depend on j and ϵ_{ij}^t is a small fluctuation; see e.g., (Bayati and Montanari, 2011, Appendix A).

To express the AMP algorithm using vector notation, we define \mathbf{m}^t and \mathbf{v}^t be $n \times 1$ vectors containing the means and variances, respectively, at iteration t . Also, we define \tilde{Y} and $\tilde{\Lambda}$ to be $n \times n$ symmetric matrices with diagonal entries equal to zero and off-diagonal entries given by \tilde{y}_{ij} and $\tilde{\lambda}_{ij}$, respectively. The AMP algorithm in (87) can be written compactly as

$$\mathbf{a}^t = \frac{1}{\sqrt{N}} \tilde{Y} \mathbf{m}^t - \frac{1}{N} (\tilde{Y}^{\circ 2} \mathbf{v}^t) \circ \mathbf{m}^{t-1} \quad (88a)$$

$$\mathbf{b}^t = \frac{1}{N} \tilde{\Lambda} (\mathbf{m}^t)^{\circ 2} + \frac{1}{N} (\tilde{\Lambda} - \tilde{Y}^{\circ 2}) \mathbf{v}^t \quad (88b)$$

$$\mathbf{m}^{t+1} = \eta(\mathbf{a}^t, \mathbf{b}^t) \quad (88c)$$

$$\mathbf{v}^{t+1} = \eta'(\mathbf{a}^t, \mathbf{b}^t) \quad (88d)$$

where η and η' are obtained by stacking the functions η_i and η'_i , respectively, for $i = 1, \dots, n$.

Remark 1. In the AMP algorithm given in (88) The dependence on distribution of x_i is encapsulated by the function $\eta_i(a, b)$ defined in (84). This means that the algorithm can be applied in the case where x_i is drawn according to a probability measure P_i provided that the functions

$$(a, b) \mapsto \int x_i^p \exp\left\{a x_i - \frac{1}{2} b x_i^2\right\} \exp\left\{\frac{\tilde{y}_{ii} x_i^2}{\sqrt{N}} - \frac{\tilde{\lambda}_{ii} x_i^4}{2N}\right\} dP_i(x_i), \quad p = 0, 1, 2 \quad (89)$$

can be approximated numerically.

Remark 2. If the parameter λ_{ii} is small relative to the global scale parameter N , then the data point \tilde{y}_{ii} has negligible impact and the term $\phi_i(x_i)$ in (84) can be replaced by the density of x_i . This is precisely what happens in the usual rank-one estimation model where all the λ_{ij} are identical. By contrast, the heteroskedastic model (67) allows for the possibility that the diagonal terms λ_{ii} are of the same order as N and in this case, the diagonal terms should not be discarded.

AMP for Groupwise Model The AMP presented in Section 4.3 is obtained by specializing the AMP algorithm in (88) to the setting of the groupwise model. Following the discussion in Remark 2 the function $\eta_i(a, b)$ defined in (84) is replaced by the definition given in (27), which omits the dependence on the diagonal entries of the observations.

C NUMERICAL APPROXIMATION OF RELATIVE ENTROPY AND MMSE

Numerically approximating the relative entropy and MMSE requires solving a low-dimensional saddle point problem and generating expressions for the single-letter relative entropy and its first derivative. In this work, we only consider numerical approximation of the formulas for the two-group case, as in (29). However, the single-letter relative entropy and its first derivative will be the same in the K -group case and our strategy for estimating the global optimum of the saddle point problem would be a reasonable strategy for $K > 2$ as well.

C.1 Estimation of the Global Optimum

In the two-group setup, we only have to perform an optimization over four scalar variables in order to obtain the relative entropy and MMSE. To further simplify the process, we can note that the envelope theorem (see Milgrom and Segal (2002)) gives us that we need only consider the stationary points of the objective. At first glance of (30), this is far from obvious. In practice, we use a multivariate root finding algorithm from the SciPy Jones et al. (01) optimization package in order to find the locations where the gradient is equal to zero.

C.2 Approximation of the Single-letter Relative Entropy

In order to deal with saddle point problem in (30), we need to obtain a formula for the single-letter relative entropy, defined in (7), and its first derivative. In this work, we consider standard Gaussian priors and discrete priors. In the case of standard Gaussian priors, we have $D_k(\gamma) = \frac{1}{2}(\gamma - \log(1 + \gamma))$ and $\partial_\gamma D_k(\gamma) = \frac{1}{2} \frac{\gamma}{1 + \gamma}$. We also consider discrete priors which have the following form

$$P_X = \sum_{i=1}^M p_i \delta_{a_i} \quad (90)$$

where δ_{a_i} denotes a point mass function at a_i and p_i are probabilities for $i = 1, \dots, M$ and P_X denotes the probability measure. In order to obtain the single-letter relative entropy, we consider the random variable $Y = \sqrt{\gamma}X + Z$, where $X \sim P_X$ and $Z \sim \mathcal{N}(0, 1)$ are independent. The probability density function of Y is a Gaussian mixture, which is given by

$$f_Y(y) = \sum_{i=1}^M p_i \mathcal{N}(y; \sqrt{\gamma}a_i, 1) \quad (91)$$

where $\mathcal{N}(y; \mu, \sigma^2)$ denotes the Gaussian probability density function with mean μ and variance σ^2 evaluated at y . The single-letter relative entropy function and its derivative can be expressed as one-dimensional Gaussian integrals, which are approximated numerically.

D ADDITIONAL FIGURES

Gaussian Priors The case where both \mathbf{u} and \mathbf{v} have independent standard Gaussian entries is shown in Figure 3. In this setting, the asymptotic MMSE is constant for all $0 \leq \alpha \leq 1$ and the weak detection threshold (i.e., the smallest value of λ such that the asymptotic MMSE is strictly less than the MMSE obtained without any observations) occurs at the critical threshold $\lambda^* = 1$. The empirical performance of both gradient descent and AMP shows excellent agreement with the formulas for the asymptotic MMSE. By contrast, the spectral methods are sensitive to the choice of the parameter α . More specifically, the spectral methods perform well at the special values of α for which the two-group model reduces to one of the previously studied spiked matrix models. (These values are given by $\alpha = 1/2$ for joint PCA and $\alpha \in \{0, 1\}$ for weighted PCA.) However, as α deviates from these special values the performance degrades sharply, indicating that these spectral methods are sub-optimal in general.

Non-Gaussian Priors Two examples involving non-Gaussian priors are shown in Figure 4. In the first example, both \mathbf{u} and \mathbf{v} have independent Rademacher entries and in the second example, \mathbf{u} has independent Bernoulli(0.1) entries (shifted and scaled to mean zero and variance one) and \mathbf{v} has independent standard Gaussian entries. In both cases, empirical results are shown only for AMP. The asymptotic MMSE is invariant to α in the first example, where \mathbf{u} and \mathbf{v} have the same distribution, but depends on α in the second example, where \mathbf{u} and \mathbf{v} have different distributions. Interestingly, the MMSE curves in the second example are ordered for both the MMSE of $\mathbf{u}\mathbf{u}^\top$ and $\mathbf{v}\mathbf{v}^\top$, however, the order is reversed between the two. In particular, smaller values of α result in better performance (lower MMSE) in estimating $\mathbf{u}\mathbf{u}^\top$, but have worse performance (higher MMSE) in estimating $\mathbf{v}\mathbf{v}^\top$.

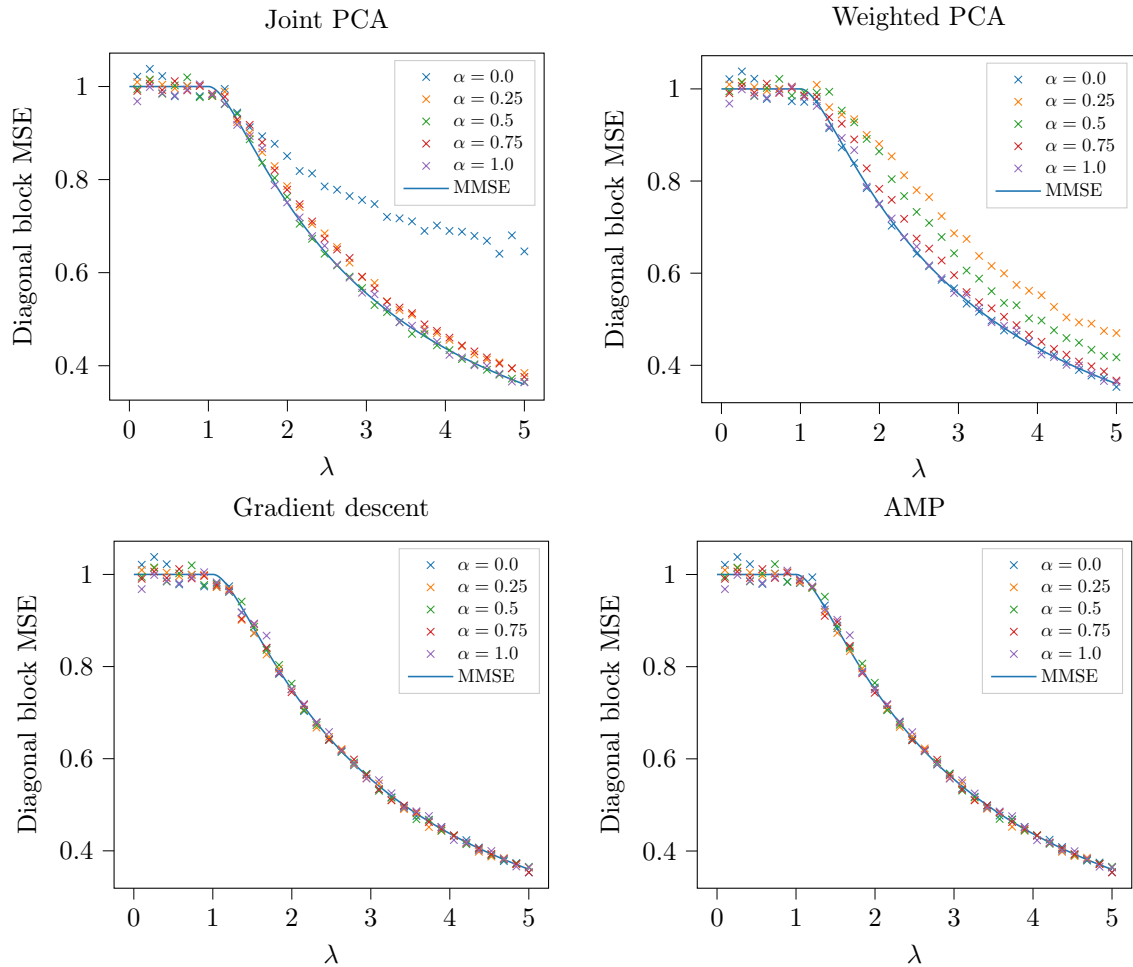


Figure 3: MSE in estimating the rank-one matrices $\mathbf{u}\mathbf{u}^\top$ and $\mathbf{v}\mathbf{v}^\top$ in the two-group model (29) as function of λ for various α with $n = 512$ and IID standard Gaussian priors. The asymptotic MMSE (solid line) is given by Theorem 2. The subplots show the diagonal block MSE of joint PCA (upper left), weighted PCA (upper right), gradient descent (lower left), and AMP (lower right).

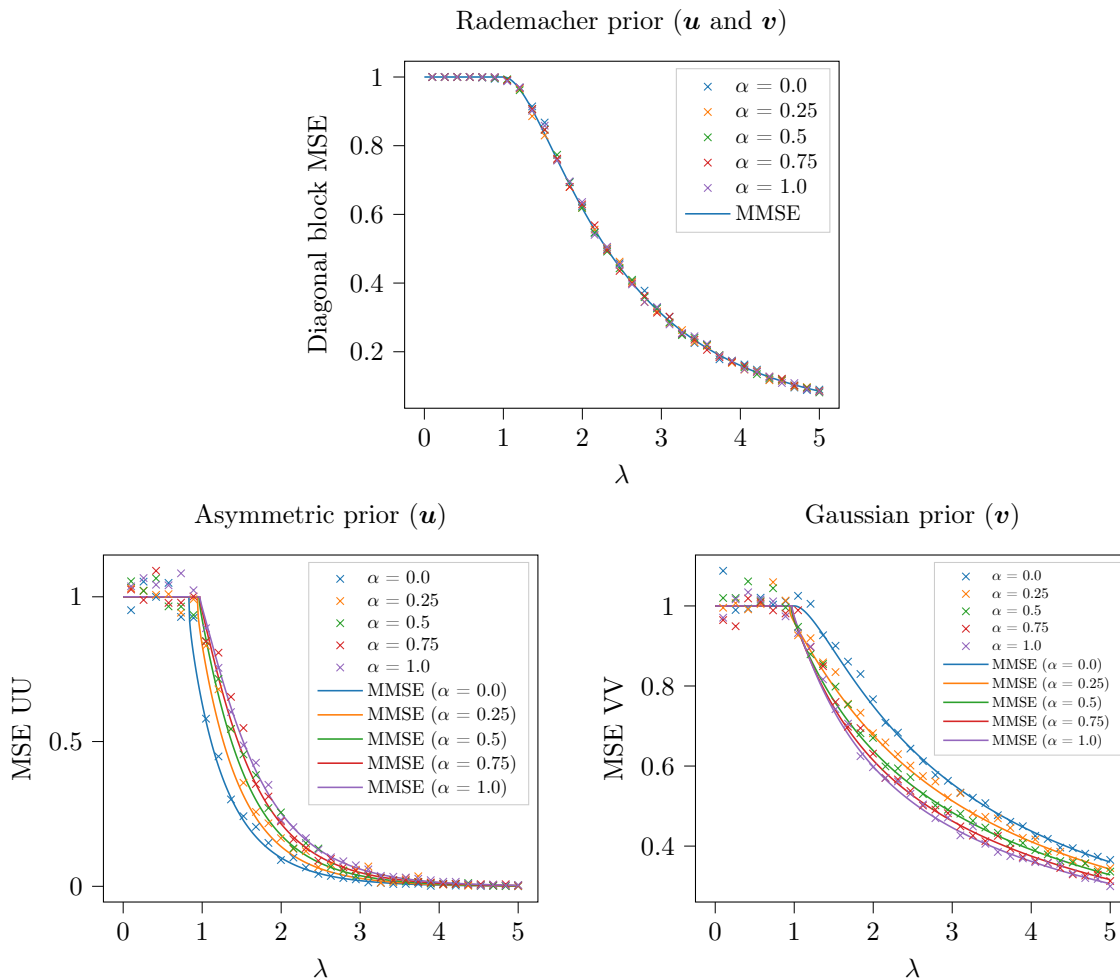


Figure 4: MSE in estimating the rank-one matrices $\mathbf{u}\mathbf{u}^\top$ and $\mathbf{v}\mathbf{v}^\top$ in the two-group model (29) as function of λ for various α with $n = 512$. The asymptotic MMSE (solid line) is given by Theorem 2. The empirical MSE corresponds to AMP and is averaged over 32 Monte Carlo trials. In the top subplot, both \mathbf{u} and \mathbf{v} are Rademacher. In the bottom subplots, \mathbf{u} has a Bernoulli(0.1) prior (scaled and shifted to zero mean and unit variance) and \mathbf{v} is standard Gaussian. In the lower left subplot, we are plotting the MMSE and MSE in estimating $\mathbf{u}\mathbf{u}^\top$ and in the lower right subplot, we are plotting the MMSE and MSE in estimating $\mathbf{v}\mathbf{v}^\top$.